ESSAYS ON SCREENING IN INFORMATION MARKETS

ELIA SARTORI

A DISSERTATION

PRESENTED TO THE FACULTY

OF PRINCETON UNIVERSITY

IN CANDIDACY FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE

BY THE DEPARTMENT OF ECONOMICS

ADVISER: STEPHEN MORRIS

JUNE 2019

Abstract

In three essays, this dissertation studies the production and distribution of information goods. In the first chapter, we model information as a digital good. Digital goods are produced along a quality ranking and can be both *duplicated* and *damaged* at zero marginal cost. Consumers' valuation of quality consists of a common decreasing returns component and an heterogeneous component that gives sellers a motive for screening. The monopolist problem is naturally divided into an acquisition and a distribution stage; two interdependent sources of inefficiency, underprovision and quality damaging, emerge. Competition is modeled as a two stage game of perfect information. Welfare comparisons between monopoly and duopoly are ambiguous: additional underacquisition and double spending favor the former, undoing damaging inefficiencies by distributing a positive quality for free favors the latter.

The second chapter studies the production of socially relevant information: we model policymaking as a bandit problem where the arms are treatment incentive schemes whose payoff value and correlation is disciplined by an economic theory. We preliminarily associate each multiarmed bandit problem to an uncertainty function so that the implied information function is traded-off one for one with expected utility at each belief state to determine the optimal policy. The uncertainty measure quantifies the estimation content of selection mechanisms. We propose a sampling procedure that validly implements all BDM mechanisms while minimizing the variance of the empirical propensity score and preserving information continuity. Fully voluntary mechanisms are control optimal under linear preferences, but their valid implementation induces the largest variance of the sample size used for estimation.

In the third chapter (with Franz Ostrizek) we study a monopolist screening problem with network externalities in consumption and two dimensions of heterogeneity: consumer differ in their susceptibility and influence (to the network effect). We show that the allocation is inefficient if and only if susceptibility is unobservable, while consumers receive rents for their influence only if susceptibility is unobserved and influence is verifiable. The optimal allocation under private information satisfies lexicographic monotonicity; bunching arises around the switching types in the lexicographic order, i.e. highest-influence types adjacent to the next level of susceptibility.

Acknowledgments:

This dissertation, as well as the academic career that I have in front of me, owes immensely to the wise advice and continuous support of my advisor Stephen Morris. Stephen shaped my experience as a graduate student by firstly convincing me to come to Princeton, then by teaching me advanced topics in information economics and finally by leading the complicated transition to becoming a researcher in this discipline. Looking back at this journey, I realize that his contribution cannot be overestimated.

I had the privilege to discuss my work with Faruk Gul and Wolfgang Pesendorfer, who I also thank for supporting my decisions. Beyond the advices on the specific projects, conversations with them helped me getting an idea of how to approach research. For example, when I start thinking about modeling some economic interaction now I find myself always asking myself "If I had infinite proving power, what result would I prove?"

The years of my PhD have been enlightened by my friends Andrei, Giorgio, Michele and Franz. In particular, I find it impressive how happy (for lack of a better word in my limited vocabulary) I feel when I am with Franz.

Al mio papà, che è difficile da ringraziare a voce.

Contents

List of Tables

List of Figures

Chapter 1

Competitive Provision of Information Goods

1.1 Introduction

We study the distribution of goods that are produced along a quality ranking and that can be both duplicated and damaged at zero marginal cost. A firm that creates a version of the good of quality q can, at no additional cost, sell arbitrary amounts of any quality below q. Consumers demand at most one version of the good, they agree on the quality ranking but have different tastes for quality. Such heterogeneity makes producers with market power willing to engage in inefficient quality damaging for screening purposes as in the literature of multiproduct monopolist as in Mussa and Rosen (1978), hereafter referred to as MR.

Two markets whose functioning is well approximated by this model are the market for digital content, (computer software, mobile apps, digital audio and video content), and some portions of the market for information (weather forecasts, non-strategic financial information). The former is a large and growing sector of advanced economies, while the latter is of interest because it is natural to assume that sellers have access to a free garbling technology that allows for damaging of information structures. Until Section 4, where we formalize the application to information markets, we will not discuss how particular results explain phenomena observed in those markets, and generically refer to digital goods as products that have the following characteristics:

1. They are non-rival but excludable through a pricing system;[1]

[1] An essential non-rivalry arises because of the free replicability. In general, no consumption externalities are allowed, which is a particularly restrictive assumption in the market for information

2. Produced and damaged along single-dimensional quality ranking that is given exogenously and on which all consumers agree. This excludes the possibility of horizontal differentiation across consumers. However,

3. Consumers have heterogenous tastes for quality.

4. When multiple firms are active, their products are homogenous: Individuals never want to combine qualities sold by different producers.

5. On the production side, replication and damaging of a version occurs at negligible cost, and damaging must lie on the pre-specified quality ranking.

We will assume consumers' preferences for quality (point 3 above) are separable in a common decreasing returns component and an heterogeneous constant returns component. We adopt this specification for two reasons. First, the interpretation we give in Section 1.2.1 in which agents use the digital good to perform two tasks (basic and professional activities) may be a reasonable description of the demand for digital goods (we indeed motivate it with an example of software consumption). Second, it is a parsimonious specification that allows for rich empirical implications of the model. The standard linear preferences used in MR, which will be presented as a subcase, are unable to generate an optimal contract that displays non-trivial damaging: we show that if the decreasing returns component were absent, then only one positive quality would be sold in the market. This is counterfactual, at least in some markets: to give just one example, Figure 1.1 below shows a set of packages for statistical software that differ in their computational power. Likewise, digital content is often offered in SD or HD packages and Section 4 gives examples of non-trivial screening in information markets.

An additional observation motivating our analysis is that goods of positive quality are distributed for free in many of the markets with the characteristics described above. An enormous amount of information is available at no (monetary) cost. The same is true for online services, ranging from e-mails to document storage and digital contents. As for computer software, it is interesting to notice that Open

where a recent literature focused on the fact that the value of information is an equilibrium object. The excludability issue is also critical; a whole literature (Muto (1986), Varian (2000), Polanski (2007) among others) focuses on the distribution of non-excludable "information goods" that in their definition include also software and books.

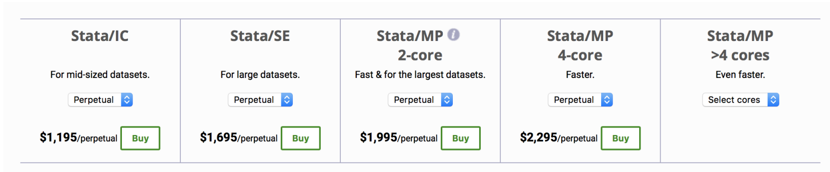

Figure 1.1: Non-trivial damaging in the distribution of Stata

Office was released in 2002, 12 years after Microsoft sold the first Office package, possibly as a consequence of increased competition. Many mobile apps are also sold for free, even though premium options are often present and, unlike (some) computer software, are immune to failures of non-excludability.[2] The literature offers some explanations for the free-quality phenomenon, ranging from creating costumer fidelization to be exploited in parallel markets (bundling and cross-fidelization), to earning profits from individual attention (through advertising). In this paper we build a simple model that can both create, under certain assumptions about primitives and the competition structure, the positive implications discussed above (non-trivial damaging and distribution of free quality), and that is flexible enough to answer some questions like: Is such free distribution socially desirable? What is the welfare impact of some policies in this framework (damaging prohibition, linear taxation, patent protection)?

1.1.1 Outline of the paper and preview of results

Section 2 formalizes the primitives and characterizes the provision of digital goods by a monopolist under both perfect and asymmetric information. As in standard first-degree price discrimination problems, the former setting induces the first-best allocation which features no quality damaging. The latter problem is more complicated. We show that it can be conveniently rewritten as the maximization over a sequence of MR problems parametrized by the quality cap constraining the mo-

[2] Many online apps profit from matching demand and supply for either transportation (Uber, Lyft) or for food delivery (Deliveroo, Foodora). They also actively engage in some sort of screening: ride hailing apps charge a higher price for larger or more comfortable cars, while food delivery apps may propose early delivery for a surcharge. Those are not examples that fit our description, since both "premium" services (bigger car and fast delivery) require a higher marginal cost to the producer: the delivery guy has to run a motorbike rather than a bike and, similarly, the premium car has a higher depreciation/fuel cost.

nopolist's allocation function. Each problem in this sequence, not the original one, can be solved applying standard monopolist screening techniques with a particular (zero) cost function. We show that the optimal contract conditional on a quality cap allocates each type the minimum between an increasing type-dependent function and the quality cap itself. A bunching threshold moves as the quality cap increases, which raises rents of high valuation types while leaving rents of lower types unchanged. This property makes it simple to characterize the marginal revenue function and hence solve the quality-acquisition problem. The two-stage nature of the monopolist problem generates two sources of inefficiency: an acquisition inefficiency similar to standard underprovision with market power, and a damaging inefficiency from asymmetric information. The two are interdependent: distribution obviously depends on the quality constraint, and incentives to acquire depend on the revenues that damaging can achieve. In particular the efficiency at the top typical of standard screening problems is limited to a distributional efficiency: a positive measure of types never receives a damaged quality but even the highest type receives a quality below what he gets in the first best. Curvature of the common component implies that at low quality levels optimal distribution features no damaging, and that in all contracts all types receive positive quality (and surplus). With linear preferences (no concave component) a "no haggling" result holds: a cap-invariant set of types is always served the undamaged quality while others are fully excluded (given $q = 0$), so marginal revenues are constant. We conclude Section 2 by analyzing the impact of a No Screening (NS) policy, namely to prohibit the seller from engaging in quality damaging for screening purposes. When binding, the policy is proved to always worsen the underacquisition inefficiency. As for its impact on damaging, two forces operate: the NS policy mechanically prevents inefficient damaging, though it may induce the complete exclusion of some low types that received positive surplus in the unconstrained monopoly contract. We find conditions under which the NS policy is welfare improving.

Section 3 studies competition in digital goods markets as the equilibrium of a two stage game of perfect information. The first stage, investment in quality, determines firms' market power at the pricing stage. The second stage (pricing) equilibrium is easily characterized using the tools developed to solve the monopolist problem: the owner of the larger quality behaves indeed as a (interim-)monopolist on the quality spectrum he owns exclusively, while Bertrand forces drive the price of the second highest quality to zero. In the first stage there are multiple equilibria indexed by n, the number of firms that are active (i.e. choosing to acquire a positive quality with positive probability). With $n = 1$ the active firm is a monopolist for sure (the only pure strategy equilibrium) and agents receive the monopolist allocation of Section 2. With any $n \geq 2$ there is a symmetric equilibrium in which active firms randomize investment with full support ranging from zero to the monopolist quality. In the class of equilibria with active competition, we prove that every type in the economy is better-off with a lower intensity of competition (smaller n). The relevant welfare comparison is therefore between monopoly and duopoly. By only looking at the support of the mixed equilibrium we observe that the highest quality distributed under competition will be below the monopolist quality (increasing underprovision inefficiency). Also, active competition implies inefficient double spending: due to homogeneity, the development cost of an inferior quality is always socially wasteful. However, by distributing a positive quality for free, competition shrinks screening inefficiencies associated to each realized best quality. We prove that we cannot go beyond this qualitative comparison between monopoly and duopoly equilibria: different shapes of the cost function can shut down almost completely either the positive or the negative impact of competition. In particular, if the monopolist was not damaging, then competition unambiguously reduces total welfare. By contrast, if costs are extremely convex (approaching a fixed cost structure), then a competitive market induces stochastic allocations that converge to the flat allocation where everybody receives the quality produced (but not distributed) by a monopolist and hence dominates the monopolistic equilibrium.

Section 4 is independent of the other two and assesses the fitness of the framework presented before to study information markets, relative to existing approaches that model information acquisition. We discuss how modeling choices used in the paper translate into implicit assumptions on the type of information markets that can be analyzed. The key element is that production of information is decentralized and that the technology to convert the factor of production (attention) into state-signal structures must be taken as a primitive. We present a simple but exact microfoundation of the reduced form model studied in the paper and discuss how correlation in primary information structures can be used to create product heterogeneity which is necessary to avoid some implications of the model that are counterfactual in those market (only one firm makes profits). The dimensionality of feasible signal structures and of payoff relevant types, even with a "small" state space, poses significant tractability challenges that are beyond the scope of this paper but suggest directions for future research.

1.1.2 Literature review

In this section we review technical contributions related to the building blocks of the model: the demand side, the composition of replicability and damaging on the production technology, and a description of model of competition with screening. Models of information markets are reviewed separately in the last section.

Quality screening

The demand side of the economy and the solution techniques for the quality-conditional problem are based on the literature on screening with a multiproduct monopolist pioneered by Mussa and Rosen (1978), advanced in Maskin and Riley (1984) and later in Wilson (1993). Assumptions in this paper make sure to avoid ironing and other technical complications within each MR problem (which are the focus of the original paper and Rochet and Choné (1998)), and to have a simple revenue comparison across different problems.

Free Replicability

Motivated by the example of software and digital contents, a recent literature in computer science Goldberg and Hartline (2003); Goldberg et al. (2006); Hartline and Roughgarden (2008) studied how to design the revenue maximizing mechanism to allocate a good that is replicable for free to agents with heterogeneous valuations. In particular Goldberg et al. (2001) show that posted price mechanism performs surprisingly close to the optimal (possibly dynamic) incentive compatible auction. This result partially justifies my focus of screening through (a menu of) prices.[3]

Damaging goods

The idea of damaging a good for screening purposes was originally introduced in Deneckere and Preston McAfee (1996).[4] The approach in this paper is different both in modeling choice and in the type of questions addressed. From a modeling perspective, beyond preserving a positive marginal cost from distribution,[5] Deneckere and McAfee (1996) take a binary set of qualities as exogenously fixed, thereby excluding an acquisition margin, and assume that the only way to produce the good of low quality is by damaging the high quality good. Marginal distribution costs are

[3]It is not a complete justification to my approach as they don't allow damaging of the replicable good. An extension of their results to cases in which the seller can damage the good would be interesting per se.

[4]Srinagesh and Bradburd (1989) offer a very general analysis for the case where there are two types of customer. McAfee (2007) provides an exact characterization in terms of marginal revenues of when damaging is profitable.

[5]Their motivating examples include processors, printers and other technological products.

therefore larger for the low quality good.[6] They focus on the monopolist problem and address the following question: When is it the case that the possibility to damage benefits all agents in the economy?[7]

Product versioning through quality damaging has been explored also in the context of a durable good monopolist. In related papers, Inderst (2008) and Hahn (2006) consider an environment with two consumer types and a monopolist that sells different versions of a product over time and faces a Coasian commitment problem in price and quality.[8]

Availability of a damaged quality may also result from illegal activities such as piracy (assuming the copied version is somehow inferior to the original).[9] Peitz and Waelbroeck (2006) provides a critical overview of the theoretical literature that addresses the economic consequences of end-user copying, though focusing mostly on the non-excludability of low qualities that is induced by the illegal activity.[10]

Competition with screening

A vast literature on competition with screening spurred from the seminal contribution of Rothschild and Stiglitz (1976) (RS) on insurance markets. RS equilibrium invokes a natural notion of stability induced by a free entry condition which can be interpreted as the Nash equilibrium of a contract[11] posting game among many (ex-ante) symmetric firms. Existence is not guaranteed. The advancement of game theory allowed a more formal analysis of the strategic interaction among competitors: Jaynes (1978) shows the RS equilibria always exist if sharing of information about customers is treated endogenously as part of the game among firms. Hellwig (1988) shows that this is true if each firm's communication strategy is conditioned on the set of contracts that are offered by the other firms.[12]

---

[6] Clearly, under those assumptions it is a more startling fact that a monopolist is sometimes willing to engage in screening.

[7] Section 1.2.5 addresses a natural extension of this question within our framework.

[8] In their setting the seller may optimally choose to engage in quality deterioration in the first period (and trade only occurs in this period) by selling the low-quality version below marginal cost in the first period to avoid later price concessions to high-valuation consumers.

[9] Takeyama (1994) argues that the loss in profits due to copying may be greater if dynamic effects are taken into account; however it is also possible that the seller benefits from being copied since this reduces his commitment problem.

[10] Hence connecting mostly with the literature on non-excludable "information goods" referred to in footnote 1.

[11] In their case, a contract is a set of insurance rate and price.

[12] With the use a dynamic games, the analysis of contracting with adverse selection was extended to the other issues, among others that of renegotiation (e.g. Hart and Tirole (1988)) and recontracting (Beaudry and Poitevin (1995) model a financial market where the informed party has the bargaining power even though competing uninformed parties make the offer).

Departures from modeling equilibrium with asymmetric information as the outcome of an extensive form game produced some elegant characterizations. Dubey and Geanakoplos (2002) study the RS model by fixing an exogenous set of pools characterized by their limits on contributions. Households signal their reliability by choosing which to join. They put discipline on beliefs over pools that are not visited in equilibrium and prove existence (and uniqueness) of the separating RS equilibrium. Bisin and Gottardi (1999) and Bisin et al. (2011) extend the the model of general competitive equilibrium to economies with asymmetric information without having to explicitly model private information.[13]

Possibly due to a lack of tractability of the latter models, the literature even in recent years has kept analyzing competition under asymmetric information as the equilibrium of an extensive form game. This is the approach taken also in this paper.[14] Netzer and Scheuer (2010) extend the RS model to two-dimensional heterogeneity in both risk and patience where the latter is the endogenous result of optimal savings or labor supply decisions. In their model RS equilibria exist and equilibrium contracts can earn strictly positive profits because any contract that attracts good consumers would also attract bad risk types and become unprofitable. A similar result of profitable contracts in RS equilibria is obtained in the two dimensional screening model of Smart (2000) where both dimensions are exogenous.

However, in many cases RS equilibria fail to exist due to the many deviations available to the pool of potential entrants. One solution in this case is to give firms some market power at the pricing stage. Garrett et al. (2014) have a model in which market power is given exogenously by assuming that also consumers are imperfectly informed about the offers in the market (two-sided asymmetric information). They show that the intensity of competition decreases this source of market power, so in the limit the Bertrand equilibrium emerges. Having firms commit through an ex-ante irreversible investment is a second way of creating (this time endogenously) market power, used since Hotelling (1929). A standard reference for this approach is Kreps and Scheinkman (1983), who have firms commit to a quantity level before Bertrand competing on the realized investments. The setup closer to that of this paper is Champsaur and Rochet (1989) who analyze a MR duopoly where each competitor

[13] Wilson (1978), Dutta and Vohra (2005) and Vohra (1999) propose an extension of the core as a positive foundation of equilibria under asymmetric information.

[14] Although a strategic foundation may sometimes be an appealing feature of the model, it adds one degree of arbitrariness: by looking at the simplest model of competition and how Cournot and Bertrand equilibria differ in their implication we understand how crucial even the specification of the action space may be. Selection of the game structure is mainly driven by the tractability of equilibria it delivers; alternative specifications are discussed in Section 1.3.4.

costlessly commits to a subset of qualities and then chooses a pricing function (paying the distribution costs at this stage). In section 1.3.4 we compare the properties of the equilibria of their models with results in this paper.

1.2 The Monopolist Problem

1.2.1 Primitives and Efficiency Benchmark

Demand

The economy is populated by a unit mass of consumers. Each consumer is characterized by a utility type $\theta \in \Theta$, where Θ is a compact subset of $\mathbb{R}_+$. $F : \Theta \to [0,1]$ strictly increasing and admitting a density is the population distribution function. Utility types describe an agent's cardinal rankings over different quality versions of the digital good. Consumers' valuation for quality are assumed to take the functional form

$$u(q,\theta) = g(q) + \theta q \tag{1.1}$$

where g is a concave function representing the relative curvature of the common component of quality ranking with respect to the type dependent one.[15] The degenerate case of $g \equiv 0$, i.e linear utility, will be an important subcase for two reasons: it is the utility specification adopted in MR among others and also it delivers an extreme version of our screening results. When non-degenerate, it is assumed that g satisfies the Inada conditions

$$\lim_{x\to 0} g'(x) = \infty, \qquad \lim_{x\to\infty} g'(x) = 0 \tag{1.2}$$

Agents also own a large amount of a numeraire good and have quasilinear preferences in this good. So the demand correspondence associated to a quality pricing function $\boldsymbol{p} : Q \to \mathbb{R}$ is given by

$$D_{\boldsymbol{p}}(\theta) = \arg\max_q u(q,\theta) - \boldsymbol{p}(q)$$

[15] As quality does not have a natural metric, we can consider a more general setting in which for two increasing function g_1, g_2

$$u(x,\theta) = g_1(x) + \theta g_2(x)$$

then define quality $q = g_2(x)$ with associated cardinal rankings

$$u(q,\theta) = g_1\left(g_2^{-1}(q)\right) + \theta q$$

The cost function over the new quality space can be redefined in a similar fashion. The qualitative results would not change, what is key is that the type independent component $g_1 \circ g_2^{-1}$ is concave or, that g_1 is "more concave" than g_2. The empirical content driving the results is that the common valuation of the quality is the additive separability in types and the fact that the function multiplying type is "less concave" than the common quality ranking.

The utility specification (1.1)-(1.2) is important to deliver some properties of the optimal contract.[16] We now justify it by offering an interpretation in the context of software consumption, and we identify the analytical properties that drive our results. For an interpretation, suppose consumers use the digital good to perform two tasks. All users perform the same basic task and measure returns to quality in the accomplishment of this task according to a common decreasing returns function; they also perform an advanced task, but they have heterogeneous constant marginal return types θ, which measure the intensity of individuals' tastes for quality in the accomplishment of such task. To substantiate the assumption, we use the software (OS) example and broadly define q as computational power. The set of basic tasks include simple calculations, text editing and other activities that are performed by everyone in essentially the same way and for which the returns to quality are very steep at the beginning, but then vanish. The advanced task is a professional activity in which each consumer specializes and that may be more or less computationally intensive. At a low level of q there is little (relative) variation in marginal utilities as everyone cares essentially about the steep improvement in the performance of the basic task, while at large q the demand for improvement is driven by the use one can make in the advanced task, which is heterogeneous.

From a technical standpoint, additive separability and concavity of g gives rise to a constant difference in the marginal utility $g'(q) + \theta$ between any two types. Yet, since g' is decreasing q, their ratio

$$\frac{g'(q) + \theta}{g'(q) + \theta'}$$

is also a decreasing function of q. The type dependent component θ within the marginal valuation $g'(q) + \theta$ becomes dominant as we climb up the quality ladder. As the marginal willingness to pay for a quality improvement for a high type relative to a low type increases in the level of quality, it becomes profitable to screen type θ from θ' only when the quality level is large enough.

Production and Sale

The digital good can be produced along a continuum of versions or qualities, where $Q = \mathbb{R}_+$ is the quality space. A producer creates a version of the good of quality q

[16] In section we present less restrictive sufficient condition to preserve the structure of the result at the cost of lower tractability.

at cost $c(q)$, assumed to be increasing and convex and measured in the same units as revenues.[17] Then he can supply an arbitrary quantity of version q as well as all versions dominated by q. Formally, they operate the following production set[18]

$$Y = \{\mathbb{I}\{q' \leq q\}, -c(q)\}_{q\in Q} \tag{1.3}$$

After producing q the seller has to quote a feasible market, that is a pricing function on the restricted domain $\boldsymbol{p} : [0, q] \to \mathbb{R}$.[19] The profit maximization problem therefore reads

$$\max_{q,\boldsymbol{p}:[0,q]\to\mathbb{R}} \int_\Theta \boldsymbol{p}(D_{\boldsymbol{p}}(\theta)) f(\theta)\,\mathrm{d}\theta - c(q)$$

First Best and Perfect Information

We begin by stating and solving the first best problem of choosing a social allocation to maximize expected social utility net of acquisition costs.

Definition 1.1. The **efficient allocation** is the function$\rho^{eff} : \Theta \to Q$ that solves

$$\max_{\rho:\Theta\to Q} S(\rho) = \int_\Theta u(\rho(\theta), \theta)\,\mathrm{d}F(\theta) - c\left(\sup_\theta \rho(\theta)\right) \tag{1.4}$$

The following proposition characterizes the efficient allocation

[17] Again, using the re-definition of the quality spectrum from footnote 15, the cost reads $c_1(q) = c\left(g_2^{-1}(q)\right)$, and the substantive assumption that $c \circ g_2^{-1}$ is convex is guaranteed since it is a composition of convex functions.

[18] Consumers are in unit measure and demand at most one version, so a supply of 1 is indeed "arbitrarily large".

[19] An equivalent restatement is to let the seller choose only a pricing function $\boldsymbol{p} : Q \to \mathbb{R} \cup \{\infty\}$ and define the cost function $\bar{c}$ on the space of pricing functions as:

$$\bar{c}(\boldsymbol{p}) = c(\sup\{q : \boldsymbol{p}(q) < \infty\})$$

In this case, the monopolist would solve

$$\max_{\boldsymbol{p}:Q\to\mathbb{R}} \int_\Theta \boldsymbol{p}(D_{\boldsymbol{p}}(\theta)) f(\theta)\,\mathrm{d}\theta - \bar{c}(\boldsymbol{p})$$

Proposition 1.1. *The efficient allocation map is given by*

$$\begin{array}{rccc} \rho^{eff} & : & \Theta & \to & Q \\ & & \rho^{eff}(\theta) & \mapsto & q^{\star} \end{array}$$

where $q^{\star}$ is the unique solution to equation

$$g'(q) + \mathbb{E}_F[\theta] = c'(q) \tag{1.5}$$

It is natural that the efficient allocation has singleton image $q^{\star}$: since each individual's utility is increasing in q and the distribution of each quality below q costs the same, it will never be socially optimal to allocate a damaged good to any type. Equation (1.5) is the first order condition of problem (1.4) after noticing that the efficient allocation is flat; sufficiency is immediate. As is also standard, we notice that a monopolist that is not subject to information frictions, namely that observes each type and can charge different prices to different costumers will induce the efficient allocation rule ρ^{eff} and extract all the surplus.

Remark 1.1. (First Degree Price Discrimination) Suppose that the individual type θ were observable to the seller. Then the monopolist would replicate the efficient allocation characterized in Proposition 1.1, and make profits $S\left(\rho^{eff}\right)$.

1.2.2 Private information

In the remainder of the paper, we assume consumers have private information about their utility type. The monopolist must rely on incentive compatible market design to allocate different qualities to different types.[20]

[20] It is assumed that the principal has no "screening device" as defined in Jaynes (2006), namely he cannot obtain additional information about valuation types so that each individual must be treated as a random draw from F (assumed known). Bergemann et al. (2015) analyze the limits of (third-degree) price discrimination induced by additional information on buyers' types and show that that information can be tailored to achieve any combination of surplus such that total surplus is below the efficient level, and producer and consumer surplus are above uniform monopoly pricing and zero, respectively.

We set up the problem as a multi-agent mechanism design problem and appeal to the revelation principle to write the monopolist problem as choosing a pair of allocation and transfer rules[21]

$$(\rho, p) : \Theta \to Q \times \mathbb{R}$$

to maximize profits under incentive compatibility and rationality constraints. Compared to standard screening problems, the key novelty is that the cost of an allocation rule ρ no longer takes the additively separable form

$$\overline{c}(\rho) = \int_{\Theta} \widetilde{c}(\rho(\theta))\,\mathrm{d}\theta$$

for some primitive cost $\widetilde{c}$. By contrast, it solely depends on one statistic of the allocation rule, namely the maximum quality, so it can be written as:[22]

$$\overline{c}(\rho) = c\left(\sup_{\theta} \rho(\theta)\right) \tag{1.6}$$

With these observations at hand, we write the monopolist problem in the following way

$$\begin{array}{lll} \max_{\rho,p:\Theta\to Q\times\mathbb{R}} & & \int_{\Theta} p(\theta)\,\mathrm{d}F(\theta) - c\left(\sup_{\theta}\rho(\theta)\right) \\ & \text{s.t.} & \\ & \text{IC} & u(\theta,\rho(\theta)) - p(\theta) \geq u(\theta,\rho(\theta')) - p(\theta') \;\; \forall\theta,\theta' \\ & \text{IR} & u(\theta,\rho(\theta)) - p(\theta) \geq 0 \; \forall\theta \end{array} \tag{1.7}$$

It is worth emphasizing that under the non additively separable cost function (1.6), we cannot solve (1.7) by piecewise maximization of an appropriately defined type dependent profit. However, the simple form of non-separability characterizing (1.6) suggests that problem (1.7) can be divided in two stages. First, revenues conditional

[21]The steps for rewriting the monopolist problem presented in the previous section as the design of a direct mechanism are standard and therefore omitted. It should not create confusion that from now on the pricing function p has domain the type space Θ rather than the quality space Q.

[22]In a companion project (joint with Franz Ostrizek) we explore screening under the more general cost structure

$$\overline{c}(\rho) = \int_{\Theta} \widetilde{c}(\rho(\theta), \rho)\,\mathrm{d}\theta$$

that allows the cost of producing the quality sold to type θ to depend on the whole allocation rule. This is useful to describe less extreme versions of economies of scale, learning, or nontrivial cost of quality replication and versioning.

on each quality cap are calculated, and then those revenues are compared with the cost of buying a quality cap. The following Lemma formalizes this intuition and introduces the revenue and cap-conditional allocation functions.

Lemma 1.1. *Define the quality constrained revenue function $V : Q \to \mathbb{R}$ given by*

$$\begin{aligned} V(q) \quad \longmapsto \quad & \max_{\rho,p:\Theta\to Q\times\mathbb{R}} \int_\Theta p(\theta)\,\mathrm{d}F(\theta) \\ & IC\ ,\ IR \\ & \rho(\theta) \leq q, \qquad \forall\theta\in\Theta \end{aligned} \tag{1.8}$$

and let $\rho_q : \Theta \to Q$ be the optimal quality allocation of problem (1.8). The solution to problem (1.7) is characterized by a quality cap q^M given by:

$$q^M = \arg\max_q V(q) - c(q) \tag{1.9}$$

and by an allocation

$$\rho^\star : \Theta \to Q = \rho_{q^M}(\theta)$$

Once we are given the V function from (1.8), it is clear we can solve (1.9) as a simple maximization in single variable (if we show V is concave, q^M will be characterized by a first order condition alone). The challenge is then to find function V and the cap-conditional allocation rule

$$\rho : Q \times \Theta \to Q$$

where $\rho(q,\theta)$ is the quality assigned to type θ when the quality cap is q.[23] This is the objective of the following section.

1.2.3 Solution of the Screening Problem

Characterizing the constraint-conditional optimal contract (and therefore the revenues) is in principle a complicated problem, since for each $q \in Q$ we need to solve for a function ρ_q. The main result of this section, Proposition 1.1, shows that whenever primitives satisfy regularity conditions, the set of constraint-conditional allocation

[23] The policy function $\rho_q : \Theta \to Q$ associated to the quality constrained problem (1.8) is the section at q of the above defined ρ; in this sense my notation is consistent and I will use $\rho(q,\cdot)$ and $\rho_q(\cdot)$ interchangeably.

rules take a simple form in which we just compare a function of the type with the constraint. Before moving to the general case with a continuous type space, we consider a simpler economy populated by only two types since it provides the intuition for the more general case and clarifies the nature of the regularity assumptions.

Two Types Example

Suppose the economy is populated by a fraction π of high marginal valuation types denoted by H and $1-\pi$ low marginal valuation types L. The following proposition characterizes the optimal contract $V(q)$, $\rho_q(\cdot)$ for this economy.

Proposition 1.2. *Let $y^\star \in Q$ be the solution to*

$$\frac{g'(y)+\theta_L}{g'(y)+\theta_H} = \pi \tag{1.10}$$

The optimal allocation takes the simple form

$$\rho(L,q) = \begin{cases} q & if\ q < y^\star \\ y^\star & else \end{cases}, \qquad \rho(H,q) = q$$

firm's profits are given by

$$V(q) = \begin{cases} u(q,L) & if\ q < y^\star \\ u(y^\star, L) + \pi\left(u(H,q) - u(H,y^\star)\right) & else \end{cases}$$

So that

$$V'(q) = \begin{cases} u'(q,L) & if\ q < y^\star \\ \pi u'(H,q) & else \end{cases} = c'(q)$$

characterizes per Lemma 1.1 the optimal quality cap.

Proposition 1.2 suggest how one may construct all cap-dependent contracts: using only the demand primitives of the model, i.e. the curvature of g and the distribution of types, we determine the point $y^\star$.[24] Then, one obtains the optimal allocation $\rho(\theta,q)$ by choosing the minimum between a type dependent threshold (in this case,

[24] As $\lim_{q\to 0} h(q) = 1$, $\lim_{q\to 0} h(q) = \frac{\theta_L}{\theta_H}$, then existence of threshold $y^\star$ requires that $\frac{\theta_L}{\theta_H}$ is not larger than π. Uniqueness follows from monotonicity of ratio of marginal utilities

$y^{\star}$ for low types, ∞ for high types) and the quality cap itself. By solving a single equation, (1.10), we can characterize the allocations for each quality cap and then infer transfer from the binding incentive constraints to construct the revenue and marginal revenue function. Finally, expression (1.10) suggests why a strictly concave g function is needed for a nontrivial allocation rule to realize. The following Remark formalizes this.

Remark 1.2. (Linear Utility). Under the linear utility specification equation (1.10) would read

$$\pi = \frac{\theta_L}{\theta_H}$$

which is not a function of q and cannot determine a threshold $y^{\star}$. In this case the monopolist would either always (i.e. at every quality cap) sell to both types, if $\theta_L > \theta_H \pi$, or always serve only high types. Notice expression (1.10) would also be trivial if we assumed any form multiplicatively separable utilities $u(q,\theta) = u_1(q) \cdot u_2(\theta)$. The fact that linear utility has no "service margin" as a function of quality will be preserved in the more general setting.

From now on we will work on a continuous type space.

Definition 1.2. We say primitives are **regular** if the utility function takes the form (1.1) and the type distribution F has a monotonically increasing hazard rate

$$h(\theta) = \frac{1 - F(\theta)}{f(\theta)}$$

Now consider the virtual valuation function

$$vv(\theta, q) = u(q,\theta) - h(\theta)\, u_\theta(q,\theta)$$

and define the correspondence of maximizers of the virtual valuation

$$\beta : \Theta \rightrightarrows Q \cup \{\infty\}$$

$$\beta(\theta) \longmapsto \arg\max vv(\theta, q)$$

where the abuse $\arg\max vv(\theta, q) = \infty$ is adopted whenever $vv(\theta, q)$ is strictly increasing in $q \in Q$.

The following Lemma, which follows from standard application of supermodular comparative statics, will be used in the proof of the main Proposition of this section.

Lemma 1.2. *If primitives are regular, then β is single valued, monotonically increasing and it is equal to ∞ on a set of positive measure $\left[\widetilde{\theta}, \overline{\theta}\right]$, where $\widetilde{\theta}$ is the unique solution to*

$$\theta - h(\theta) = 0$$

The general expression of β is given by

$$\beta(\theta) = \begin{cases} (g')^{-1}(h(\theta) - \theta) & \theta < \widetilde{\theta} \\ \infty & \theta \geq \widetilde{\theta} \end{cases}$$

Since g' is decreasing by assumption and so is $\theta - h(\theta)$, then β is an increasing continuous function that asymptotes to ∞ in the interior of Θ. Also, by the Inada condition and $h(\underline{\theta}) - \underline{\theta} > 0$ we obtain $\beta(0) > 0$, which will deliver important implications for the shape of the cap-contingent optimal contract. The graph below plots an example of β.

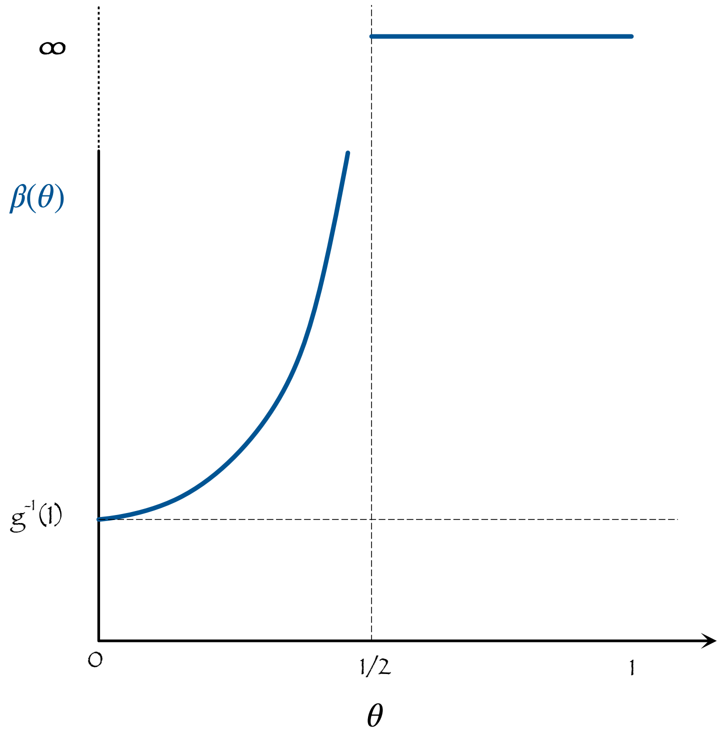

Figure 1.2: β function for concave preferences.

The threshold $\widetilde{\theta}$ takes value $\frac{1}{2}$ in the plot since in the remainder of this section we will assume $\boldsymbol{\theta} \sim \mathcal{U}[0,1]$. As Section formalizes, qualitative results are unchanged though explicit formulas for revenues, surplus and inefficiencies would need to carry an additional transformation of the inverse hazard rate.[25] Under the uniform distribution assumption, $h(\theta) - \theta = 1 - 2\theta$, $\beta(0) = (g')^{-1}(1) > 0$, and $\widetilde{\theta} = \frac{1}{2}$. We now proceed to state the proposition characterizing the monopolist allocation.

Theorem 1.1. *Suppose primitives are regular and types are uniformly distributed. Then,*

i) The ***quality-contingent*** *optimal contract takes the simple form*

$$\rho(q,\theta) = \min\{q, \beta(\theta)\} \tag{1.11}$$

[25]We do not have a clear interpretation for the distribution of types, as their empirical content is not independent of the valuation function (1.1).

ii) The revenue function V is concave with ***continuously differentiable*** *derivative given by*

$$V'(q) = \begin{cases} g'(q) & g'(q) > 1 \\ \left(\frac{1+g'(q)}{2}\right)^2 & g'(q) \leq 1 \end{cases} \tag{1.12}$$

iii) The monopolist quantity q^M is always strictly below $q^\star$.

The allocation rule (1.11) proves that the intuition in the two type example extends to regular continuous type spaces: the cap-contingent allocation of each type is simply determined by comparing a type dependent function (found at an ex-ante, constraint-free stage from Lemma 1.2), with the quality cap itself. It is only convenient to screen type θ if the quality cap exceeds $\beta(\theta)$, and when this happens the allocation $\rho(q, \theta)$ becomes unresponsive to further cap increments.

The shape of the optimal contract(s) helps interpreting also expression (1.12), i.e., the marginal revenue function. To understand this result, and for future reference, it is convenient to define $b : Q \to \left[0, \frac{1}{2}\right]$ as the inverse β function that returns for each quality level the lowest type that is bunched at the top. Under the uniform distribution we get

$$b(q) = \beta^{-1}(q) = \max\left\{0, \frac{1 - g'(q)}{2}\right\} \tag{1.13}$$

At an intermediate step of the derivation of the marginal revenue we get

$$V'(q) = (1 - b(q))\left[g'(q) + b(q)\right]$$

This expression has an intuitive interpretation: the term $g'(q) + b(q)$ is the marginal utility of the "marginally bunched" type $b(q)$, while $(1 - b(q))$ is the mass of types above him.[26] By marginally increasing the quality cap, the monopolist does not alter the revenues made from optimal allocation of all lower qualities (given to the same types at the same price). He allocates the marginal quality to type $b(q)$ and to all those above, who increase their marginal transfer by

$$u_q(q, b(q)) = g'(q) + b(q)$$

Point *iii)* states that the top quality distributed by a monopolist is below the efficient level implied by (1.5). Although not obvious in this setting, the result

[26]Notice that when $b(q) = 0$, then $V'(q) = g'(q)$ delivering the first branch in (1.12). Substitution of the nontrivial expression for b delivers the second branch.

suggests a natural parallelism with the classic underprovision of a good in the presence of market power. The "efficiency at the top" result, which is typical of quality screening problems, does not hold in this framework. Despite the highest type never receiving a damaged quality,[27] the quality cap produced under monopoly is below the level $q^{\star}$ that type, and everyone else, receives in the first best.

Comparing *i)* and *iii)* in Theorem 1.1 with Proposition 1.1 we notice that a monopolist induces two sources of inefficiency: one from damaging, because generically, $\rho(q,\theta) < q$ for a set o positive measure, and one from suboptimal acquisition. Although associated to different stages of the monopolist problem, these inefficiencies are to some degree interdependent as (1) the screening allocation is clearly constrained by the quality acquired and (2) the quality acquired depends on the screening possibilities. Expression (1.12) incorporates the optimal distribution of each maximal quality, which generally entails damaging. A graphic representation of the inefficiencies is given in the bottom panel of Figure 1.3, where green represents underacquisition and orange damaging. Their analytical expression will be derived in the next section.

1.2.4 Properties of the Monopolist Contract

We now list some properties of the optimal contract which follow immediately from Theorem (1.1). We begin with a description of the resulting quality allocations, separating the case of strictly concave and trivial g.

Corollary 1.1. *(Linear utility) Suppose $g \equiv 0$. Then, optimal allocations are given by:*

$$\rho(q,\theta) = \begin{cases} 0 & \theta \leq \frac{1}{2} \\ q & \theta > \frac{1}{2} \end{cases}$$

and the monopolist has constant marginal revenue $V'(q) = \frac{1}{4}$.

[27] So that distributional efficiency at the top realizes in our setting.

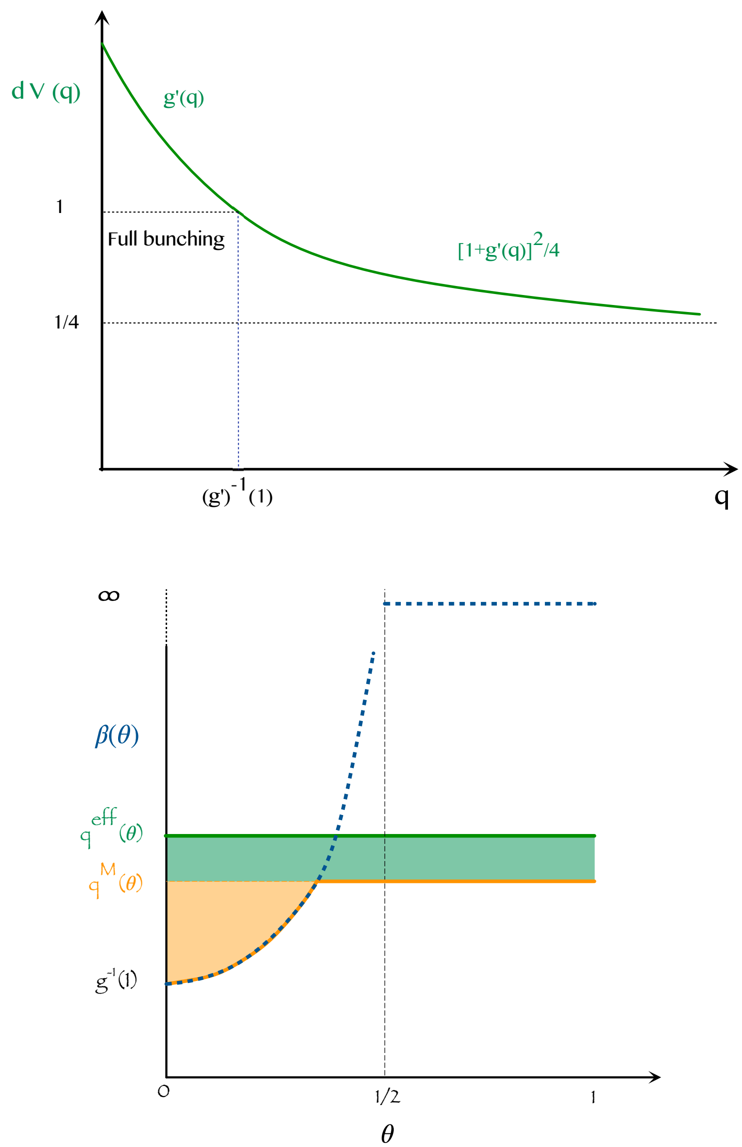

Figure 1.3: Marginal revenue (top) and the two inefficiencies (bottom) from monopolistic provision of digital goods.

The result follows by applying allocation rule (1.11) to the maximizer of the virtual valuation, which in case of linear preferences is:[28]

$$\beta(\theta) = \begin{cases} 0 & \theta < \frac{1}{2} \\ \infty & \theta > \frac{1}{2} \end{cases}$$

When preferences are linear, screening is only performed by excluding (selling quality 0 to) an invariant set of types, so that only one positive quality will be offered. All the non-trivial screening observed in price discrimination with linear preferences is driven by the shape of the marginal cost curve (and from ironing a non-monotone hazard rate). Since the distribution problem lacks such cost curvature, a smooth screening contract must result from the specification of preferences. Allocations with concave g have the following properties

Corollary 1.2. *Suppose g is strictly concave. Then*

i) All types receive a positive quality in the optimal contract.

ii) At low quality caps the optimal contract features full bunching:[29]

$$\rho(q,\theta) = q \quad \forall \theta, q < (g')^{-1}(1)$$

iii) A positive measure of agents $\left[\frac{1}{2},1\right]$ receives the highest quality good irrespectively of the quality cap.[30]

$$\rho(q,\theta) = q \quad \forall \theta \in \left[\frac{1}{2},1\right]$$

Point *i)* is implied by the Inada condition of g around 0: by giving a marginal quality to low valuation types who receive nothing the seller gets unbounded marginal

[28] As in Remark 1.2, the shape of the optimal contract would be the same for any multiplicatively separable utility specification

$$u(q,\theta) = u_1(q) \cdot u_2(\theta)$$

which gives:

$$vv(q,\theta) = u_1(q) \cdot \left[u_2(\theta) - u_2'(\theta)\left(\frac{1-F(\theta)}{f(\theta)}\right)\right]$$

and $\beta(\theta)$ is again either 0 or ∞.

[29] With non-uniform distribution, the requirement is $q < (g')^{-1}(h(\underline{\theta}) - \underline{\theta})$

[30] Similarly, we would get the unconditional bunching region to be $\left[\widetilde{\theta}, \overline{\theta}\right]$.

revenues, which she can distribute as information rents to make sure IC constraints for higher types are satisfied. If the quality acquired is low enough, since there is little variation in relative marginal utilities it will not be optimal to screen any type, giving point *ii*). An implication of *ii*) is that sufficiently steep marginal cost shuts down one source of inefficiency, damaging, leaving only underacquisition active. This will allow us to isolate the impact of a policy on the underacquisition inefficiency alone by assuming monopolist was producing in this region. By point *iii*), there is a set of types that are bunched at the top irrespectively of the quality cap. It should be noticed that those types are exactly the same that were sold a positive undamaged quality under linear preferences. Indeed, a concave g does not change the fact that above a certain threshold the virtual valuation is monotonically increasing, but it gives a nontrivial maximizer for other types. Also, we can now explain why (1.12) gives limit marginal revenues

$$\lim_{q \to \infty} V'(q) = \frac{1}{4}$$

As the quality grows higher, marginal increments are distributed as if preferences were linear since at high qualities agents use incremental units only in the performance of the advanced task.[31]

Figure 1.4 shows a graphical derivation of the optimal contract: the marginal revenue is crossed with the marginal cost function to determine q^M (left graphs with axes flipped for convenience), this level is reported on the vertical axis in the graph of β and the optimal allocation $\rho^\star(\cdot)$ then results from "slicing" β at q^M. The three panels describe full bunching, active screening and linear preferences.

We now compute total surplus under monopoly and give an analytic expression to the two sources of inefficiency described above.[32] If q^M is below $(g')^{-1}(1)$, then computations are simple. Producer revenues are $V(q) = g(q)$, consumer surplus is $W(q) = \frac{1}{2}q$ and per Corollary 1.2 *i*), there are no damaging inefficiencies (only underacquisition is active). The following proposition characterizes surpluses when instead $q^M > (g')^{-1}(1)$.

Proposition 1.3. *(Decomposition of inefficiencies above $(g')^{-1}(1)$)*

[31] This also means that to have a finite solution to the monopolist problem the marginal cost must have limit that exceeds $\frac{1}{4}$.

[32] We are currently missing an expression for consumers' rent.

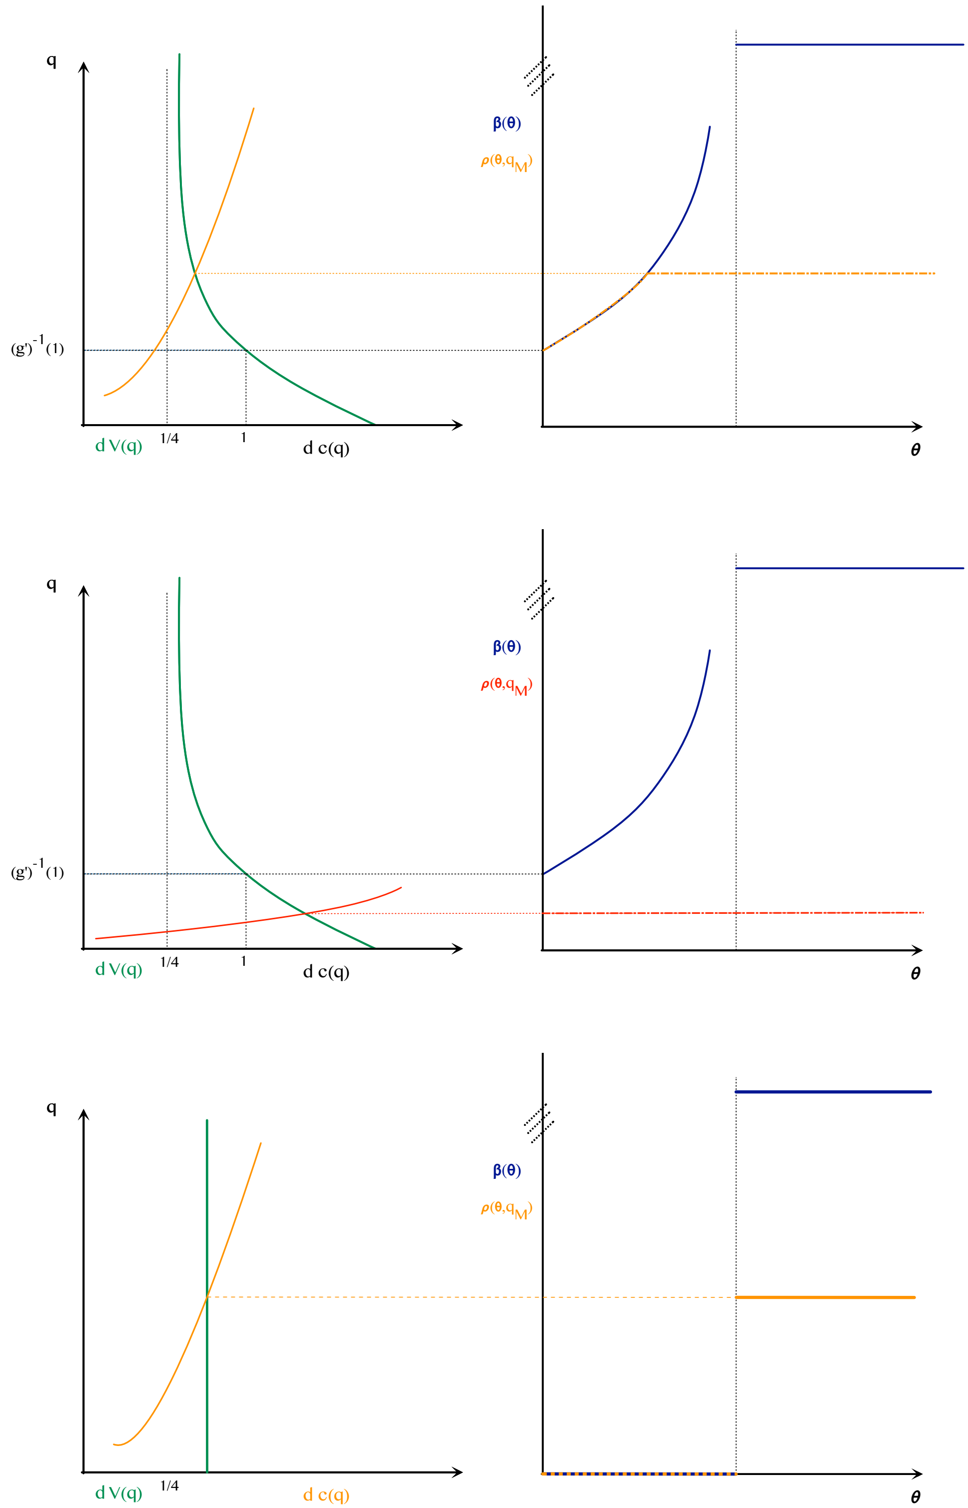

Figure 1.4: Graphical representation of optimal contract with high (top) and low (medium) marginal cost for concave g, and for linear preferences (bottom).

Total consumer surplus is given by

$$W\left(q^M\right) = \frac{1}{2}\left[\left(g'\right)^{-1}(1) + \int_{(g')^{-1}(1)}^{q^M} \left(\frac{1+g'(q)}{2}\right)^2 \mathrm{d}q\right] \tag{1.14}$$

We can therefore decompose monopolist inefficiencies as

$$\int_{(g')^{-1}(1)}^{q^M} d(q)\,\mathrm{d}q + \int_{q^M}^{q^\star} \left[\frac{1}{2} + g'(q) - c'(q)\right] \mathrm{d}q \tag{1.15}$$

where

$$d(q) = \frac{1}{8}\left(1 + \left(2 - 3g'(q)\right) g'(q)\right) \tag{1.16}$$

are the marginal inefficiencies from damaging.

Analytic manipulation, shows that aggregate marginal information rent is given by

$$W'(q) = \frac{1}{2}\left(\frac{1+g'(q)}{2}\right)^2 = \frac{1}{2}V'(q)$$

Hence, total surplus under monopoly grows, above $(g')^{-1}(1)$ with slope given by

$$\underbrace{\left(\frac{1+g'(q)}{2}\right)^2 - c'(q)}_{\text{marginal profits}} + \underbrace{\frac{1}{2}\left(\frac{1+g'(q)}{2}\right)^2}_{\text{marginal rents}} = \frac{3}{2}\left(\frac{1+g'(q)}{2}\right)^2 - c'(q)$$

Which, combined with full bunching below $(g')^{-1}(1)$ gives the following expression for marginal monopolist surplus

$$m(q) = \begin{cases} g'(q) + \frac{1}{2} - c'(q) & g'(q) > 1 \\ \frac{3}{2}\left(\frac{1+g'(q)}{2}\right)^2 - c'(q) & g'(q) \leq 1 \end{cases} \tag{1.17}$$

If we didn't have the inefficient damaging, total surplus would grow with slope

$$\frac{1}{2} + g'(q) - c'(q) \tag{1.18}$$

Subtracting the two we get the expression for the marginal inefficiencies from damaging

$$d(q) = \frac{1}{8}\left(1 + \left(2 - 3g'(q)\right) g'(q)\right)$$

by integrating $d(q)$ from $(g')^{-1}(1)$ to q^M we get the first term in (1.15) (the area of the orange region in Figure 1.3). While we cannot give an intuitive interpretation as to why (1.16) represents marginal inefficiencies from damaging, because the analytical result depends on the uniformity assumption, we notice that $d(q) = 0$ when $g'(q) = 1$ (below that level there were no damaging inefficiencies), and that d is a positive hump shaped function in $g'(q)$ when smaller than 1. Underacquisition inefficiencies, the green area constitute instead the second summand in (1.15) and add to social surplus losses (1.18) in the underacquisition region $\left[q^M, q^\star\right]$.

Relaxing Demand Primitives

We have made restrictive assumptions about the specification of returns from quality and of the distribution of types. Returns belong to the family (1.1) and the type distribution is assumed uniform. The latter assumption is innocuous: the qualitative results of allocation rule (1.11), monopolist underprovision and decomposition of inefficiencies into damaging and underprovision always hold though the analytic expression of marginal revenues (1.12), consumer surplus and damaging inefficiencies (1.16) are modified to allow for a different inverse hazard rate. Under a general distribution characterized by hazard rate h, the full bunching threshold $(g')^{-1}(1)$ would be

$$\overline{q} = (g')^{-1}\left(\underline{\theta} - h\left(\underline{\theta}\right)\right)$$

where $\overline{q} > 0$ is guaranteed by the monotone hazard rate assumption and the Inada condition $\lim_{x\to 0} g'(x) = \infty$. Similarly, the inverse β function used to calculate the marginal type bunched at the top which gives marginal revenues and welfare is defined implicitly by the equation:

$$g'(q) = (h(b(q)) - b(q)) = \widetilde{h}(b(q))$$

In general, the b function would be given by

$$b(\cdot) = \max\left\{0, \widetilde{h}^{-1} \circ g'(\cdot)\right\}$$

which reduces to (1.13) in the uniform case as$\widetilde{h}^{-1}(x) = {}^{(1-x)}/_{2}$.The cutoff type assigned the undamaged quality when preferences are linear will be $\widetilde{\theta}$, the zero of

$\widetilde{h}(\theta) = \theta - h(\theta)$,[33] and marginal revenues would be $\widetilde{\theta}\left[1 - F\left(\widetilde{\theta}\right)\right]$, which are $\frac{1}{2}$ and $\frac{1}{4}$ respectively in the uniform case.

The additive separability assumption is more substantial; its empirical content is discussed in the introduction, together with a plausible microfoundation. The Inada conditions are essential to get the full bunching region at low qualities, while identifying preferences with function g allows to characterize revenues and welfare only as a function of its curvature. For the bunching at the top property of the cap-contingent contract (Proposition 1.1, *i*)) it would be sufficient that β is an increasing function. This would be guaranteed by concavity in q and supermodularity of the virtual valuation which is in turns ensured by the standard conditions on mixed derivatives

$$u_{qq} < 0, \qquad u_{q\theta} > 0, \qquad u_{q\theta\theta} \leq 0$$

and a monotone hazard rate.

1.2.5 Impact of a No-Screening policy

Corollaries 1.2 *ii*) and 1.1 present two cases in which only one positive quality is distributed under monopoly: either preferences are linear, a constant mass of agents is served the undamaged quality and others are excluded, or g is concave but the optimal quality cap is low enough (steep marginal cost) to induce full bunching. The aim of this section is to evaluate the positive and normative implications of a regulation that prohibits the monopolist from selling damaged goods. This exercise is useful for two reasons. First, it is the natural extension to this framework of the Deneckere and McAfee (1996) normative question "when is the possibility of screening beneficial for all types in the economy?". Our different specification of the cost function and the fact that available qualities are not pre-determined add different channels through which the NS policy can impact allocations and welfare. Second, this is a first pass at evaluating the impact of a policy on the two inefficiency sources isolated in the previous section, and will provide a useful benchmark of comparison for the welfare implications of competition.

The NS Problem

The superscript NS will denote objects associated to a No Screening monopolist. He chooses a quality q^{NS} and the threshold consumer $\vartheta\left(q^{NS}\right)$ who is indifferent

[33]That always coincides with the threshold type that is always bunched at the top even with concave concave g.

between purchasing the good and not, then sell q^{NS} to types $\left[\vartheta\left(q^{NS}\right),1\right]$ at price

$$g\left(q^{NS}\right)+\vartheta\left(q^{NS}\right)q^{NS}$$

The exclusion policy $\vartheta : Q \to \Theta$ is found by solving the quality-conditional pricing problem:[34]

$$\Pi^{NS}\left(q\right)=\max_{\theta}\left[u\left(q,\theta\right)\right]\left(1-\theta\right)$$

At the acquisition stage, the NS monopolist solves

$$\max_{q}\Pi^{NS}\left(q\right)-c\left(q\right)$$

The following proposition characterizes the solution in regular environments. We focus on strictly concave g as we already argued that the NS policy is not binding when preferences are linear.

Proposition 1.4. *i) The exclusion policy ϑ is given by*

$$\vartheta\left(q\right)=\max\left\{\frac{q-g\left(q\right)}{2q},0\right\}$$

It holds $\vartheta\left(q\right)\leq b\left(q\right)$, strictly when $b\left(q\right)>0$. Moreover, $\vartheta^{-1}\left(0\right)>\beta\left(0\right)$.

ii) Marginal revenues for the NS monopolist are given by

$$\frac{\mathrm{d}}{\mathrm{d}q}\Pi^{NS}\left(q\right)=\begin{cases}g'\left(q\right) & q\leq g\left(q\right)\\ \frac{\left(q+g\left(q\right)\right)\left(q-g\left(q\right)+2qg'\left(q\right)\right)}{4q^{2}} & q>g\left(q\right)\end{cases}$$

iii) It holds $q^{NS}\leq q^{M}$ (strictly whenever $q^{M}>\left(g'\right)^{-1}\left(1\right)$).

The ϑ and b functions are plotted on the left panel of Figure 1.5. The fact that ϑ has a larger intercept, stated above as $\vartheta^{-1}\left(0\right)>\beta\left(0\right)$, is key as it implies that the NS monopolist starts excluding some types at a quality level at which the unconstrained monopolist is already actively engaging in inefficient damaging. The quality space is then partitioned in three regions, highlighted in the right panel of Figure 1.5, where marginal revenues for the constrained monopolist are compared with those of the unconstrained monopolist. In Region A, where $q<\left(g'\right)^{-1}\left(1\right)$, both monopolists do

[34]We keep the assumption $\boldsymbol{\theta}\sim\mathcal{U}\left[0,1\right]$. The qualitative properties do not change under generic monotone hazard rate distribution.

full bunching, the constraint is immaterial, and marginal revenue functions coincide. In Region B, where $q < (g')^{-1}(1)$ but $q < g(q)$, the NS monopolist sells to all types (marginal revenue is $g'(q)$) while the unconstrained one does positive damaging. In Region C also the NS monopolist performs his "constrained" screening by excluding a positive mass of low types.

Point *iii*) immediately implies that the NS policy induces a deterioration in the acquisition inefficiency, which becomes strict as soon as the NS constraint becomes binding. As for the screening inefficiencies, we have two competing forces: as $\vartheta(q) < b(q)$, conditional on acquiring the same quality, the NS monopolist would serve it to a larger portion of types. However, types that are excluded in the NS contract receive nothing, while in the unconstrained monopolist contract they have positive consumption (and value). This delivers the immediate welfare implication that the NS policy makes some low types in Region C, i.e. those that are excluded, worse off. It is also clear that the monopolist is always worse off since she is solving a constrained version of problem (1.7). We can write foregone profits from the NS policy as

$$V\left(q^{NS}\right) - \Pi\left(q^{NS}\right) + \int_{q^{NS}}^{q^M} \left[V'(q) - c'(q)\right] \mathrm{d}q$$

both summands are positive and they separate losses for not performing screening and from acquiring a quality that is below the unconstrained optimum. By focusing on cases where the unconstrained monopolist produces in Region B we can derive some less trivial welfare implications of the NS policy.

Proposition 1.5. ***(Welfare impact in Region B)*** *If the monopolist produces in Region B, then*

i) *A set of (low) types is better-off under the NS policy.*

ii) *The net gains from enacting the NS policy can be expressed as*

$$\int_{(g')^{-1}(1)}^{q^{NS}} d(q)\,\mathrm{d}q - \int_{q^{NS}}^{q^M} m(q)\,\mathrm{d}q \tag{1.19}$$

where $m(q)$ is the marginal monopolist surplus (1.17) and $d(q)$ are the marginal damaging inefficiencies (1.16).

iii) *If $c''\left(q^M\right)$ is large enough (approaches the fixed cost limit), the NS policy increases total welfare (net gains approach $\int_{(g')^{-1}(1)}^{q^M} d(q)\,\mathrm{d}q$).*

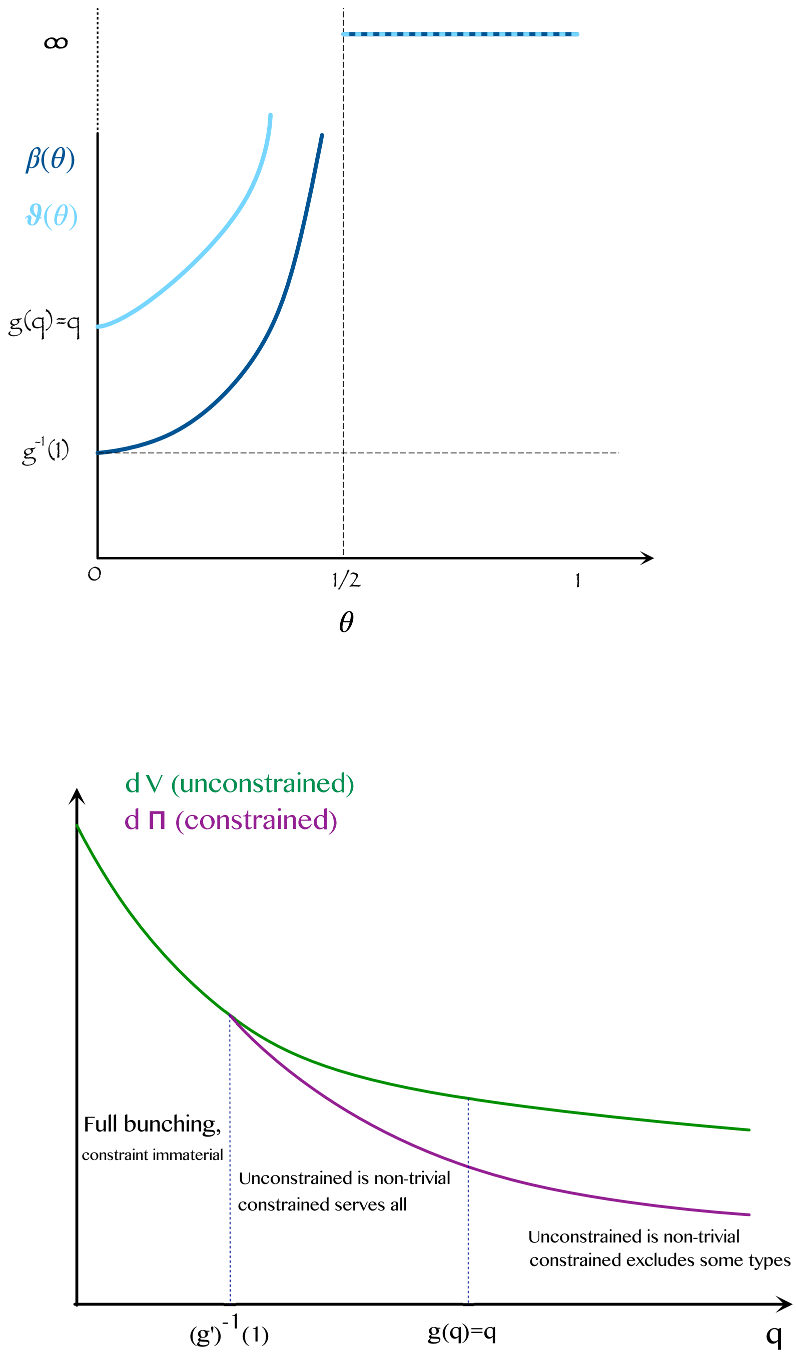

Figure 1.5: Constrained and unconstrained monopolist: full bunching thresholds (top) and marginal revenues (bottom).

The intuition for i) is as follows: because in Region B there is no exclusion, every type receives welfare $W^{NS}(\theta) = \theta q^{NS}$; in the unconstrained case low types, those below $b\left(q^M\right)$, received $W(\theta) = \int_0^\theta \beta(\theta')\,\mathrm{d}\theta'$. As $q^{NS} > \beta(0)$, it follows the marginal surplus is larger in the NS environment for a set of positive measure, implying the statement. As for the overall impact of the NS policy, it reduces "by brute force" the damaging inefficiencies but it creates two perverse effects: it makes the monopolist worsen underacquisition, and it may force some types to be completely excluded. Focusing on Region B ensures the "complete exclusion" margin in non-existent, so point ii) only trades off the positive impact from undoing damaging in the $\left[(g')^{-1}(1), q^{NS}\right]$ region, and the negative underacquisition impact. Assuming costs are extremely convex around q^M ensures marginal cost cover the $V' - \Pi'$ gap quickly implying $q^{NS} \to q^M$. So in the limit every type receives the quality a monopolist produces (but does not distribute), which increases total welfare by completely undoing damaging inefficiencies without perverse effects. The geometric intuition for the result is given in top panel of Figure 1.6, comparing acquisition and distribution in Region B for different cost functions.

In Region C, portrayed in the bottom panel of Figure 1.6, the convex cost limit would give an ambiguous welfare impact as we would need to take into account the loss of consumer surplus in the region $\left[0, \vartheta\left(q^M\right)\right]$: a full welfare comparison needs to take into account of consumer surplus that is lost by inducing complete exclusion in that Region.

1.3 Competition

This section develops a model of competition in digital goods markets. Following the discussion in the introduction, we specify an extensive form game and study its equilibria.[35] The discussion on stability in Section 1.3.4 offers a comparison with the equilibrium notion in Rothschild and Stiglitz (1976), which is a standard benchmark in the literature of competition with screening. Before that, we present the competition game, solve for its equilibria by backwards induction (pricing and investment stages), and analyze their welfare properties.

[35] How the setup of the game relates to existing models of competition with screening won't be discussed throughout the exposition, relevant references are given in the dedicated paragraph of the literature review. Section 1.3.4 compares some properties of equilibria we obtain with those of models that are closer to ours.

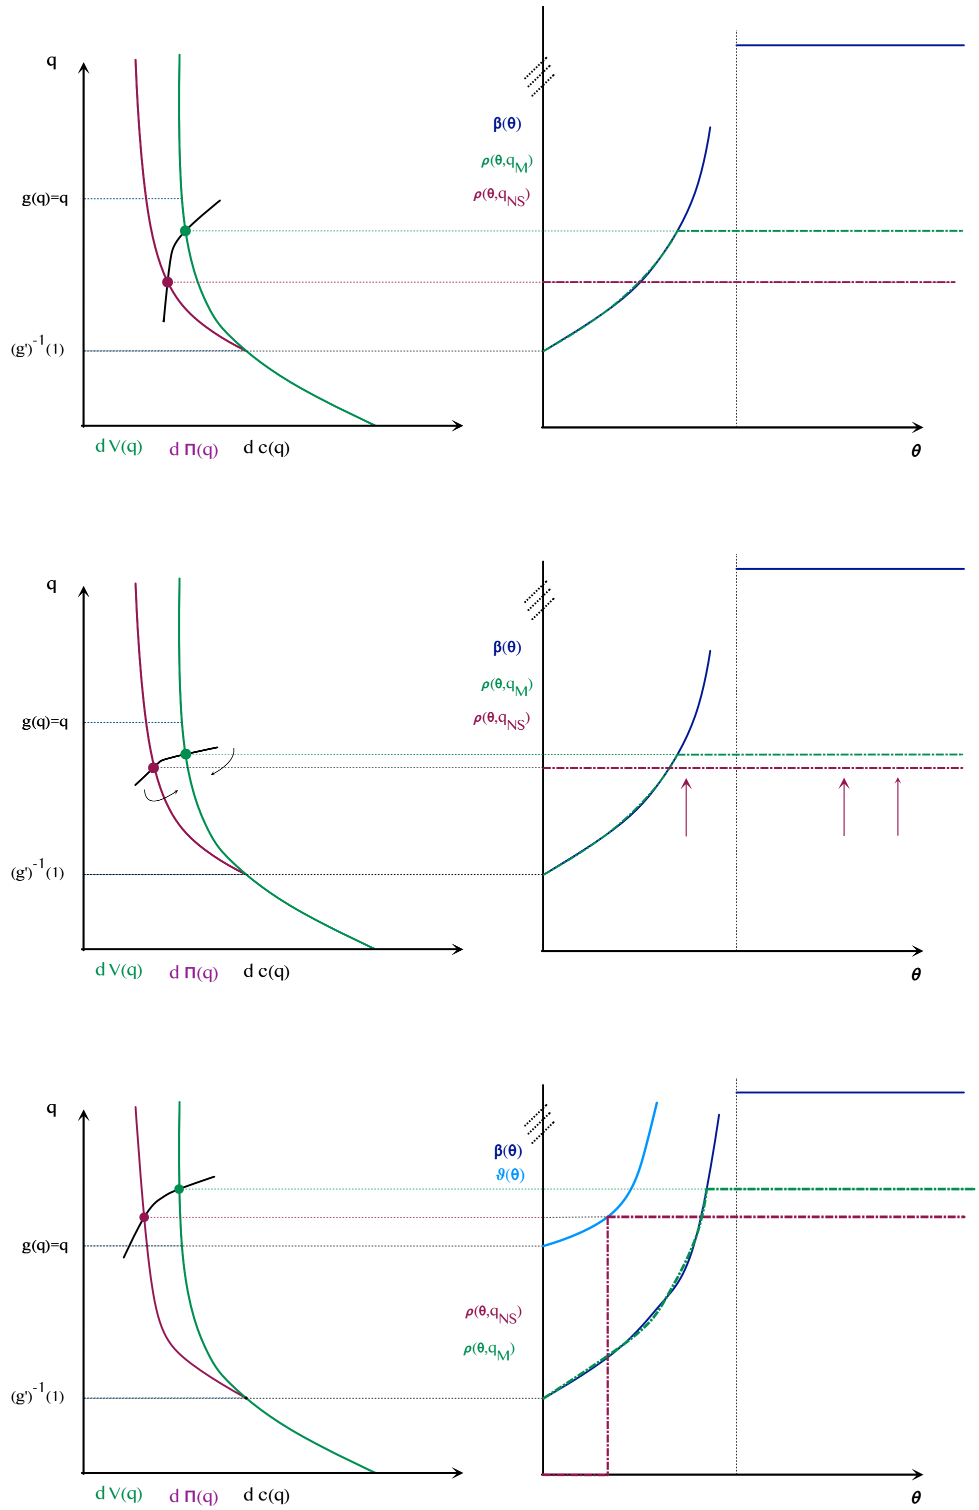

Figure 1.6: (Top panel) Unconstrained and NS allocations in Region B, cost moderately convex. (Medium) Region B, extremely convex cost. (Bottom) Unconstrained and NS allocations in Region C.

1.3.1 Primitives and the Extensive Form Game

The production primitives of the model are augmented by adding countably infinite replica of the producer studied in Section 1.2. The set of firms is denoted $\mathbb{N}$ with typical element i. On the demand side, we assume that firms produce a homogeneous (range of) products: the identity of the firm producing a certain quality is immaterial to consumers that will therefore construct their demand by looking at the lower envelope of the pricing functions. Firms play a two stage game of perfect information in which investments in quality (first stage, or acquisition stage) are based upon the belief that the ensuing price decision will constitute a Nash equilibrium in the second (or pricing) stage at which firms quote a feasible pricing function at treat production costs as already been sunk.

Timing and Action Space

At the first stage each firm chooses a mixed strategy over the qualities she acquires, so her action space is

$$A_{1,i} = \Delta\left(Q\right), \forall i$$

Firms pay the cost associated to the realized q in exchange to the right to sell (at the second stage) all qualities below q. At the end of stage one, the vector of realized qualities[36] $\boldsymbol{q} \in Q^{\mathbb{N}}$ becomes common knowledge across competitors. Conditioning on this information, firms then quote a market over the qualities available to them to maximize revenues. The action set

$$A_{i,2}\left(\boldsymbol{q}\right) = \mathbb{R}^{[0,\boldsymbol{q}_i]}$$

is the set of pricing functions on the feasible domain $[0, \boldsymbol{q}_i]$, $\boldsymbol{q}_i$ being the i^{th} entry of vector $\boldsymbol{q}$.[37] Payoffs from this stage are the revenues each firm makes as a consequence of everyone's pricing decision. A formal expression of the payoff function is given in Section .

Firms' strategies in the extensive form game consist then of a distribution over entry qualities and, conditional on each realized quality vector, a feasible pricing

[36]The bold notation reflects the fact that, from an ex-ante perspective, $\boldsymbol{q}$ is a random vector with distribution parametrized by equilibrium actions.

[37]In principle we can allow for randomization even at the pricing stage so that $A_{i,2}\left(\boldsymbol{q}\right) = \Delta\left(\mathbb{R}^{[0,\boldsymbol{q}_i]}\right)$; randomization will never be optimal in the second stage, so we restrict the action set to pure actions.

function. It is assumed that there is no discounting between periods[38] so total payoffs simply add costs incurred in the first stage and revenues earned in the second stage; when those are stochastic, firms behave as expected profit maximizers.

1.3.2 Equilibrium

This section studies subgame perfect equilibria of the finite horizon game described above by backwards induction.

The Pricing Stage

The first step consists in specifying revenues as a function of players' actions, i.e. pricing functions. Intuitively, firms make revenues from the set of qualities they offer at the lowest price, individual demands being determined by the market pricing function. The following steps make this intuition formal.

We take as given the set of individual pricing functions $\{p_i\}_{i\in\mathbb{N}}$, $p_i : Q \to \mathbb{R} \cup \{\infty\}$.[39] The market pricing function returns the lower envelope

$$m\left(\{p_i\}_{i\in\mathbb{N}}\right) : Q \to \mathbb{R} \cup \{\infty\}$$

$$m\left(\{p_i\}_{i\in\mathbb{N}}\right)(q) \mapsto \min_i p_i(q)$$

which allows to derive the individual demand correspondence

$$D_{\{p_i\}_{i\in\mathbb{N}}}(\theta) = \arg\max_q u(q,\theta) - m\left(\{p_i\}_{i\in I}\right)(q)$$

To express revenues, we firstly associate each type to the firm he buys from

$$\iota(\theta) : \Theta \to \mathbb{N}$$

$$\iota(\theta) \mapsto \min\left\{\arg\min_i \left\{p_i\left(D_{\{p_i\}_{i\in I}}(\theta)\right)\right\}\right\}$$

Firm j earns revenues that depend on how much consumers demand, which is a function of the whole market, on the quality spectrum it ends up supplying by

[38] As costs are incurred in the first period and profits are possibly earned in the second period. Re-normalizing costs to take discounting into account does not modify the analysis in any substantive way.

[39] As before, we identify firm i not offering quality q by writing $p_i(q) = \infty$.

charging the lowest price:[40]

$$\overline{R}_j\left(\{p_i\}_{i\in I}\right) = \int_{\{\theta:\iota(\theta)=j\}} p_j\left(D_{\{p_i\}_{i\in I}}(\theta)\right) f(\theta)\,\mathrm{d}\theta \tag{1.20}$$

Now suppose $\boldsymbol{q}$ is the vector of realized qualities, and that it is common knowledge across players. As it is standard, the notation $\boldsymbol{q}^{(i)}$ denotes the i^{th} order statistic of vector $\boldsymbol{q}$. The following proposition characterizes the equilibrium of the pricing game.

Proposition 1.6. *If preferences are regular, for each* $\boldsymbol{q}$ *the second stage game has an essentially unique Nash equilibrium in pure strategies. The induced allocations are*

$$\rho(\boldsymbol{q},\theta) = \begin{cases} \boldsymbol{q}^{(2)} & \text{if } \beta(\theta) < \boldsymbol{q}^{(2)} \\ \beta(\theta) & \text{if } \boldsymbol{q}^{(2)} \leq \beta(\theta) < \boldsymbol{q}^{(1)} \\ \boldsymbol{q}^{(1)} & \text{if } \beta(\theta) \geq \boldsymbol{q}^{(1)} \end{cases} \tag{1.21}$$

Revenues are given by:

$$R_i(\boldsymbol{q}) = \max\left\{V(\boldsymbol{q}_i) - \max_{j\neq i} V\left(\boldsymbol{q}_j\right), 0\right\} \tag{1.22}$$

The intuition for the result is fairly simple: since at the pricing stage costs are sunk and the production realization is common knowledge, Bertrand competition will drive to zero the revenues from versions $\left[0,\boldsymbol{q}^{(2)}\right]$ that can be provided by more than one firm. Therefore, all firms make zero revenues, except for the owner of the highest quality, from now on referred to as "interim monopolist". She enjoys market power on the quality spectrum $\left[\boldsymbol{q}^{(2)},\boldsymbol{q}^{(1)}\right]$, and behaves as a monopolist under the additional constraint that all agents must receive at least $\boldsymbol{q}^{(2)}$ for free. Given regularity, the solution to this problem is again simple: the β function is now "sliced" both from below and from above: all types below $b\left(\boldsymbol{q}^{(2)}\right)$ receive $\boldsymbol{q}^{(2)}$ for free, the others get the same allocation as under monopolist with quality $\boldsymbol{q}^{(1)}$ - but pay less. Figure 1.7 plots the allocation induced by competition with $\boldsymbol{q}^{(1)} = x$, $\boldsymbol{q}^{(2)} = y$. The equilibrium is only essentially unique because it is not pinned down who between the producer of the first and the second quality (or both of them) ends up distributing $\boldsymbol{q}^{(2)}$.[41]

[40] ι is measurable under the assumption that firms quote an increasing function.

[41] This fact may have an empirical content, since we may observe multiple providers of the free quality, though it clearly has no implications on welfare.

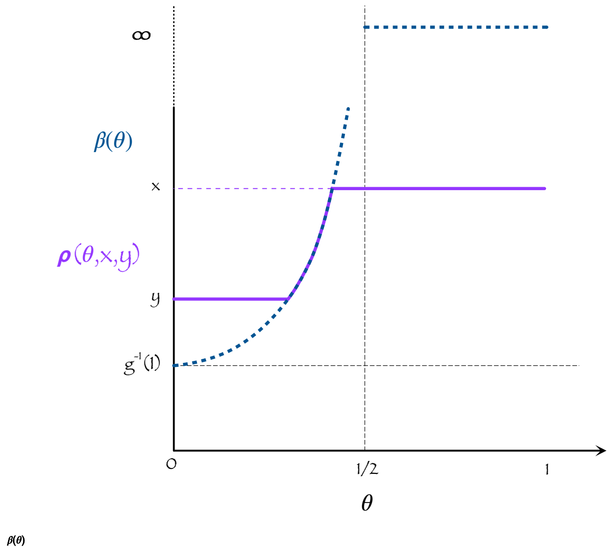

Figure 1.7: Equilibrium allocations in an x, y market.

By (1.22), each firm can compute its revenues as a function of its quality q and the best quality across competitors x, which a payoff sufficient summary of the competitive environment

$$R : Q^2 \to \mathbb{R}$$

$$R(q, x) \mapsto \max\{V(q) - V(x), 0\} \tag{1.23}$$

Towards the calculation of the first stage equilibrium it will be key that the marginal revenues of the interim monopolist only depend on her, and that they own quality and coincide with that of the unconstrained monopolist, namely

$$\frac{\partial}{\partial q} R(q, x) = \begin{cases} 0 & x > q \\ V'(q) & x \leq q \end{cases} \tag{1.24}$$

In order to avoid carrying order statistics notation, in the remainder of the paper we will denote with $\boldsymbol{x}$ be the best and $\boldsymbol{y}$ the second quality in the vector of realized entries; x, y will denote particular market realizations and $\rho(x, y, \theta)$ is the quality assigned to type θ whenever the realized first and second qualities are x, y respectively. The following example begins computation of the equilibrium in the case of linear preferences.

Example 1.1. With linear preferences, the allocation and transfers associated to each pair x, y are given by

$$(\rho, p)(x, y, \theta) = \begin{cases} (y, 0) & \theta \in \left[0, \frac{1}{2}\right] \\ \left(x, \frac{1}{2}(x-y)\right) & \theta \in \left[\frac{1}{2}, 1\right] \end{cases} \tag{1.25}$$

Low types, who were previously excluded, receive the second quality y for free, while high types receive the best quality x and pay price $\frac{1}{2}(x-y)$. The revenues of the interim monopolist are $\frac{1}{4}(x-y)$, and marginal revenues are constant at $\frac{1}{4}$.

The First Stage Game

The pricing game delivers a revenue function (1.23) which is added to the acquisition cost, thus determining the payoff function in the quality investment game. The expected profit associated to quality q is

$$\Pi(q) = \int_Q R(q, x) \, \mathrm{d}H(x) - c(q) \tag{1.26}$$

where H is the CDF of the best quality produced in equilibrium by competitors. The following proposition characterizes equilibria of this game. A firm is called active if it plays an action different from δ_0, that is if it chooses a positive quality with positive probability.

Proposition 1.7. *The first stage game has a unique equilibrium for any number $n \geq 1$ number of active firms. With $n = 1$, the active firm plays δ_{q^M}; this is also the only equilibrium in pure strategies. For each $n \geq 2$, active firms play a mixed symmetric equilibrium*

i) with support $\left[0, q^M\right]$

ii) and continuously differentiable (on the interior of the support) CDF H_n given by

$$H_n(q) = \left[\frac{c'(q)}{V'(q)}\right]^{\frac{1}{n-1}} \tag{1.27}$$

and make zero (expected) profits.

The intuition for the monopolist being the only investment equilibrium in pure strategies is the following. Two firms cannot commit to a positive quality as the owner of the lower quality would profitably deviate by abstaining. All firms abstaining cannot be an equilibrium as well, since everyone would best respond by playing monopolist. So one firm playing monopoly is the only candidate equilibrium in pure strategies. It is indeed an equilibrium: potential entrants by (1.23) do not want to choose a quality below q^M; by deviating above q^M a firm will be interim monopolist and make profits

$$\begin{aligned} \Pi(q) &= V(q) - V\left(q^M\right) - c(q) \\ &= \int_{q^M}^{q} V'(q') - c'(q')\,\mathrm{d}q' - c\left(q^M\right) \end{aligned}$$

Both summands are negative as as $c'(q) > V'(q)$ above q^M. This contrasts with the possibilities of an ex-post deviator in the spirit of Rothschild and Stiglitz (1976), who upon entry can make revenues approximately close to those of an "idle" monopolist.

For equilibria with active competition, i.e. $n \geq 2$, standard arguments from war of attrition games prove that each firm must play an atomless distribution over qualities, and that they must make zero profits. Using (1.24), the flat profit condition $\Pi'(q) = 0$, necessary for indifference, yields

$$H(q) = \frac{c'(q)}{V'(q)} \tag{1.28}$$

Equation (1.28) pins down the distribution of the maximal quality across competitors. Notice that $H(q) = 1$ right at $q = q^M$, so the support of the maximum across competitors, hence that of each firm, is $\left[0, q^M\right]$. Such a support restriction is implied by the fact that each firm's best response to a realized entry vector belongs to the doubleton set $\left\{0, q^M\right\}$: per (1.22) opponents' revenues are ex-post equivalent to a fixed cost, and only affect the decision of whether to enter, but not the quality upon entry.

Symmetry then implies the formula for H_n; notice the number of active firms is not pinned down. The CDF (1.27) is differentiable since V'' is continuous per Proposition 1.1 *ii*). To check that this is indeed an equilibrium, notice active firms are indifferent by construction on $\left[0, q^M\right]$, meaning that potential entrants would make expected losses by playing in that range (compete against n rather than $n-1$

players). Deviations above q^M are again excluded as profits would be

$$\Pi(q) = \Pi\left(q^M\right) + \int_{q^M}^{q} V'(q') - c'(q')\,\mathrm{d}q'$$

The first summand is zero in expectation by the flat profit condition (negative for a potential entrant), while the second term is negative by the definition of q^M. Again, the intuition is that irreversible investment gives incumbents the commitment to fight and drive expected profits to zero on the common support, so the interim monopolist with quality above q^M adds negative marginal profits to zero.

Combined with Proposition 1.6, the support restriction delivers the following

Corollary 1.3. *Irrespective of the number of active firms, with probability 1 a competitive market distributes*

i) a best quality strictly below q^M, and

ii) a strictly positive quality for free

Point *ii*) gives a simple empirical implication of the model: a positive quality is distributed for free if and only if there is active competition. If that is the case, we also know from *i*) that high valuation types receive a quality that is below their second-best allocation. Example 1.1 is expounded upon by computing first stage equilibria for linear preferences under a class of convex cost functions.

Example. (1.1 continued). Consider the class of convex cost functions $c(q) = q^\alpha$, $\alpha > 1$. Constant marginal revenues $V'(q) = \frac{1}{4}$ imply monopolist quality is

$$q^M(\alpha) = \left(\frac{1}{4\alpha}\right)^{\frac{1}{\alpha-1}}$$

Using Proposition 1.7 *i*), each one of n active firms plays in equilibrium the mixed strategy characterized by CDF

$$H_{n,\alpha}(q) = \mathbb{I}\left\{q \in \left[0, \left(\frac{1}{4\alpha}\right)^{\frac{1}{\alpha-1}}\right]\right\} \cdot (4\alpha q)^{\frac{\alpha-1}{n-1}}$$

1.3.3 Welfare

We now study ex-ante (expected) welfare across different equilibria. One active firm is the monopolist benchmark to which are associated the underacquisition and damaging inefficiencies isolated in expression (1.15). The positive implications highlighted in Corollary 1.3 give a first idea of the impact that active competition has on welfare (relative to the monopolist equilibrium). Point *i*) implies that the underacquisition inefficiency will be worsened. Conditional on each realized best quality x two additional forces operate. One is a multiple spending inefficiency: all costs associated to qualities that realize below x are social waste as a planner could achieve the same allocation at no additional cost. This operates in the same direction as the underprovision inefficiency, favoring monopolist. Point *ii*) of Corollary 1.3 however implies that competition shrinks the image of the allocation function, thus reducing damaging inefficiencies (1.16).

Qualitatively, therefore, comparison between monopolist and active competition equilibria is inconclusive: underacquisition and multiple spending favor the former, undoing distributional inefficiencies favor the latter. Also, notice that the forces into play are similar to those of a NS policy: perturbation of the monopolist environment induces a worsening of the underacquisition inefficiency, though it may have positive distributional effects. A third channel however distinguishes the welfare impact of the NS policy from that of active competition: in the former case we have the "complete exclusion" margin, while in the latter the double spending inefficiency. Also, all competition outcomes are stochastic as firms play mixed equilibria in the first stage.

The next proposition shows that the relative strength of the potential welfare impacts is not unambiguously signed. We show by means of example that depending on the shape of the cost function one can favor either monopoly or competition with two active firms, which uniformly (i.e. type-wise) dominates equilibria with more intense competition. We initially define the surplus for type θ conditional on realized market statistics x, y

$$W(\theta, x, y) = g(y) + \int_0^\theta \max\{y, \min\{x, \beta(\theta')\}\} \, \mathrm{d}\theta' \tag{1.29}$$

and expected welfare of type θ in the n-equilibrium

$$W_n(\theta) := \mathbb{E}_n[W(\theta, \boldsymbol{x}, \boldsymbol{y})]$$

where expectation $\mathbb{E}_n$ integrates market statistics under the distribution induced by the equilibrium (1.27) with n active firms. Since all producers make zero expected profits, total surplus under competition is just

$$W_n = \mathbb{E}_\theta \left[\mathbb{E}_n \left[W(\theta)\right]\right]$$

Theorem 1.2.

i) Equilibria with active competition are Pareto-ranked, decreasing in n*. Moreover, improvement is uniform in types, that is*

$$W_n(\theta) \geq W_m(\theta)$$

for all θ *and* $2 \leq n \leq m$*.*

ii) If $q^M \leq (g')^{-1}(1)$ *(monopolist does full bunching), then competition reduces welfare. Otherwise, we can specify a cost function under which double spending and underprovision inefficiencies vanish and the complete undoing of damaging inefficiencies make duopoly dominate.*

We notice that type-dependent welfare (1.29) is an increasing function in both x and y. This is a natural consequence of the fact that larger x gives extra surplus to high types leaving unaffected (allocation and rent of) low types, while larger y increases the allocation of low types and reduces payment for high. For i) it is therefore sufficient to show that the joint distribution of $(\boldsymbol{x}, \boldsymbol{y})$ is ordered according to first order stochastic dominance (FOSD) in n. We prove that the distribution of $\boldsymbol{y}$ conditional on $\boldsymbol{x} = x$ is independent of n for each x, with (conditional) distribution given by

$$H_{x,n}(y) = \frac{H(y)}{H(x)} \mathbb{I}\{y \in [0, x]\} \tag{1.30}$$

and that the distribution of $\boldsymbol{x}$ is ranked in n according to FOSD, from which sufficiency follows. By point i) we can therefore compare 2 active firms[42] with the monopoly equilibrium.

Point ii) proves that the qualitative forces highlighted in Corollary 1.3 can have any relative strength and, depending on the shape of the cost function can favor either monopoly or duopoly. The message is delivered by considering two extreme cases:

[42] We call two active firms duopoly though the term may be misleading, as inactive firms are also key players in equilibrium.

steep cost around zero and the fixed cost limit. We begin with the first case As there is no damaging, monopolist surplus is

$$W_1 = g\left(q^M\right) + \frac{1}{2}q^M - c\left(q^M\right)$$

Also, by Proposition 1.6, as $y \leq x \leq q^M$ for any market realization everyone will be allocated x at price $g(x) - g(y)$. Per (1.29), type dependent surplus is

$$W(x, y, \theta) = g(x) + \theta x - (g(x) - g(y)) = g(y) + \theta x$$

and total surplus is

$$W(x, y) = \int_\Theta W(x, y, \theta)\, \mathrm{d}\theta = g(y) + \frac{1}{2}x$$

By using the expression for the conditional distribution (1.30) and then integrating by parts, we show that welfare under competition $W_n = \mathbb{E}_n\left[W(\boldsymbol{x}, \boldsymbol{y})\right]$ can be written as

$$\underbrace{\left[g\left(q^M\right) - c\left(q^M\right) + \frac{1}{2}q^M\right]}_{W_M} - \int_0^q \left[\underbrace{g'(x) - \frac{\mathrm{d}}{\mathrm{d}x}\left(c(x)\frac{g'(x)}{c'(x)}\right)}_{>0} + \frac{1}{2}\right] H_n(x)\, \mathrm{d}x \quad (1.31)$$

First term is monopolist welfare, second term is the integral of a positive function, so the statement is proved.

We prove that in general, i.e. allowing for $q^M > (g')^{-1}(1)$, the difference between duopoly and monopolist welfare can be written as

$$W_2 - W_M = \mathbb{E}_2\left[\underbrace{\int_{(g')^{-1}(1)}^{\boldsymbol{y}} d(q)\, \mathrm{d}q}_{\text{undo screening}} - \underbrace{c(\boldsymbol{y})}_{\text{double spend}}\right] - \mathbb{E}_2\left[\underbrace{\int_{\boldsymbol{x}}^{q^M} m(q)\, \mathrm{d}q}_{\text{underacquisition}}\right] \quad (1.32)$$

where $\mathbb{E}_2$ is distribution of market statistics under duopoly equilibrium, m is marginal monopolist surplus function (1.17), d is marginal damaging inefficiencies (1.16), and the following convention is used: for any real valued function f,

$$\int_a^b f(x)\, \mathrm{d}x = 0 \quad (1.33)$$

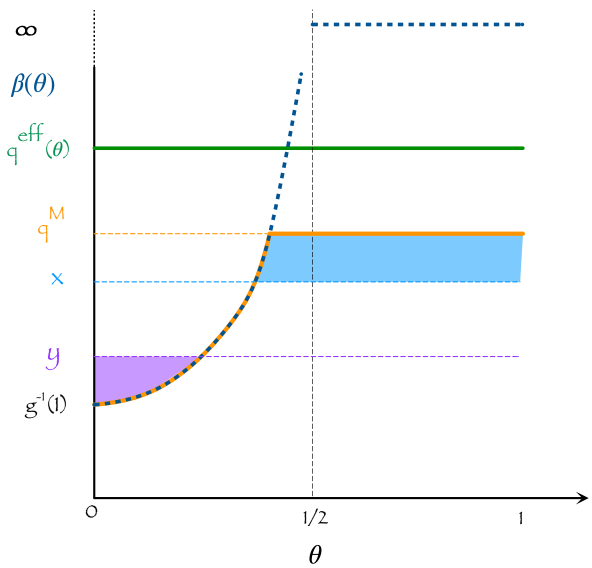

Figure 1.8: Screening undoing (purple) and underprovision (light-blue) welfare impact

whenever $a > b$. As we highlighted in the expression, each summand in (1.32) isolates the impact on one inefficiency. Figure 1.8 below plots for market realizations x, y the resulting welfare impact from undoing screening (purple), and worsened underacquisition (light-blue).

The only positive summand in (1.32) comes from realizations of $\boldsymbol{y}$ above $(g')^{-1}(1)$: unless it shrinks the allocation function compared to an $x-$monopolist, y has no impact on total welfare as it only transfers surplus from the seller to (all) consumers in the form of lower price. The case $q^M < (g')^{-1}(1)$ ensures this happens with probability one.[43] To get a positive impact of competition the following conditions are required. Both $\boldsymbol{x}$ and $\boldsymbol{y}$ should put significant mass on high realizations: the former must be close to q^M to reduce the integration domain of the underprovision inefficiency, while the latter should be well above $(g')^{-1}(1)$ to have significant undoing of damaging inefficiencies; finally, the cost of the second quality must be small to reduce the double spending inefficiency. Notice that high realizations of $\boldsymbol{y}$ have an ambiguous effect on welfare, as they increase both the integration domain for screening inefficiencies and costs.

Expression (1.32) is easily modified to evaluate welfare impact of higher intensity competition: we would need to take expectations under the a different (FOSD domi-

[43] Indeed the implication that monopolist dominates duopoly whenever the former does not actively damage could be immediately derived from (1.32). We used the more direct welfare calculations that the subcase allowed to perform.

nated) distribution of market statistics and to account for all multiple spending as all realizations below the best quality will increase the amount of wasteful acquisition.

We are left to show that in case monopoly has active damaging inefficiencies, then competition can dominate. To this end, suppose $1 > (g')^{-1}(1)$ and consider the limit of convex cost functions $c(q) = q^{\alpha}$ as α grows to infinity.[44] Irrespectively of the revenue function monopolist quality will converge to the point at which the cost function explodes, while its cost will converge to zero

$$q_{\infty}^{M} = \lim_{\alpha \to \infty} q^{M}(\alpha) = 1, \quad c_{\infty}^{M} = \lim_{\alpha \to \infty} \left(q^{M}(\alpha)\right)^{\alpha} = 0$$

also, substituting marginal cost into (1.27) we get that equilibrium strategy converges in probability to q_{∞}^{M}, that is

$$H_2(q) \to \begin{cases} 0 & q < q_{\infty}^{M} \\ 1 & q = q_{\infty}^{M} \end{cases}$$

and $\mathbb{E}_{\alpha}[c(\boldsymbol{y})] \leq c\left(q_{\alpha}^{M}\right) \to 0$. Plugging those results in (1.32) we observe that the limit welfare impact of competition is given by

$$\begin{aligned} &\int_{(g')^{-1}(1)}^{q_{\infty}^{M}} d(q)\,\mathrm{d}q - c_{\infty}^{M} - \int_{q_{\infty}^{M}}^{q_{\infty}^{M}} \left[m(q) - c'(q)\right]\mathrm{d}q \\ &= \int_{(g')^{-1}(1)}^{q_{\infty}^{M}} d(q)\,\mathrm{d}q \end{aligned} \tag{1.34}$$

only the impact on screening undoing is active in the limit, and is also “complete”: all types receive in the limit the quality that a monopolist would have produced (but not distributed).

We provided an expression for the welfare gains under duopoly and proved it cannot be unambiguously signed. Competition induces positive distributional effects that contrast increased underacquisition and double spending. The cost function can be tailored to shut down either channel. This is the main message of this Section.

We now complete the analysis of equilibria with linear preferences and generic costs started in Example 1.1 by studying their welfare properties and show how results derived in this section apply to a tractable example. Notice that with linear

[44]It is assumed $g'(1) \leq 1$ so the asymptotic monopolist engages in inefficient damaging; otherwise we can target any asymptotic monopoly level $\overline{q}$ by letting $c_{\alpha}(q) = \left(\frac{q}{\overline{q}}\right)^{\alpha}$.

preferences there is no region in which the screening inefficiency is non-existent, as types $\left[0,\frac{1}{2}\right]$ are always (completely) excluded.

Example. (Example 1.1 continued) Using allocation rule (1.25), we can write market type-dependent welfare as

$$W(\theta,x,y)=\begin{cases}\theta y & \theta<\frac{1}{2}\\ \theta x-\frac{1}{2}(x-y) & \theta>\frac{1}{2}\end{cases}$$

So we can derive a closed form expression for market-contingent welfare[45]

$$W(x,y)=\int_0^{\frac{1}{2}}\theta y\mathrm{d}\theta+\int_{\frac{1}{2}}^1\left[x\theta-(x-y)\frac{1}{2}\right]\mathrm{d}\theta=\frac{1}{8}[x+3y] \tag{1.35}$$

With quadratic cost $c(q)=\frac{1}{2}q^2$, the monopolist produces $q^M=\frac{1}{4}$, while each firm in equilibrium with n active firms plays distribution

$$H_n(q)=\mathbb{I}\left\{q\in\left[0,\frac{1}{4}\right]\right\}\cdot(4q)^{\frac{1}{n-1}}$$

From which we can calculate expectations of first and second order statistics

$$\mathbb{E}_n[\boldsymbol{x}]=\int_0^{\frac{1}{4}}\left[1-(4x)^{\frac{n}{n-1}}\right]\mathrm{d}x=\frac{n}{8n-4} \tag{1.36}$$

$$\mathbb{E}_n[\boldsymbol{y}]=\int_0^{\frac{1}{4}}4ny\left[1-[4y]^{\frac{1}{n-1}}\right]\mathrm{d}y=\frac{n}{16n-8} \tag{1.37}$$

Figure 1.9 plots equilibrium support and mean allocations (sufficient for welfare under linearity) for different competition intensities.

Plugging (1.36) and (1.37) into (1.35) we get that welfare in an equilibrium with n active firms is

$$W_n=\mathbb{E}_n[W(\boldsymbol{x},\boldsymbol{y})]=\frac{1}{8}\mathbb{E}_n[\boldsymbol{x}+3\boldsymbol{y}]=\frac{5}{64}\frac{n}{2n-1}$$

[45]Notice y has disproportionate impact as it increases welfare of low types and reduces payments of others.

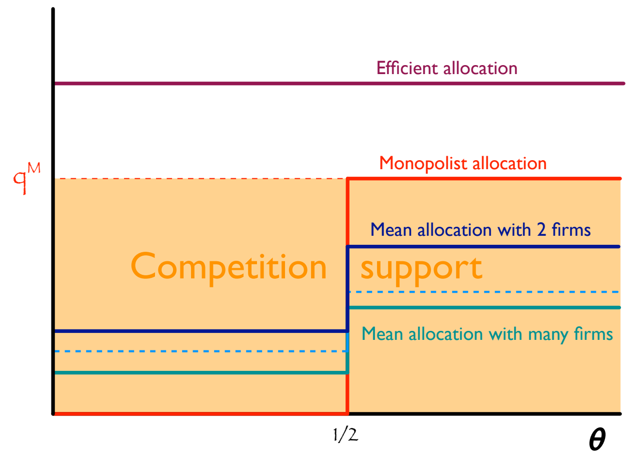

Figure 1.9: Expected Bertrand allocations with different intensity, quadratic cost.

which is a decreasing function in n (Proposition 1.7 i)). Under monopolist, both consumer and producer surpluses are $\frac{1}{32}$, giving total surplus $\frac{1}{16}$. Notice

$$W_1 = \frac{1}{16} > \frac{5}{64} \cdot \frac{2}{3} = W_2 > W_3 > \cdots > \lim_{n\to\infty} W_n = \frac{5}{128} > \frac{1}{32} = CS^M$$

With moderately convex costs monopolist outperforms duopoly and competition of any intensity makes consumers better off than under monopoly.

Now consider generic convex cost $c(q) = q^\alpha$, and focus on the limit case $\alpha \to \infty$. We can check that

$$q^M(\alpha) = \left(\frac{1}{4\alpha}\right)^{\frac{1}{\alpha-1}} \to 1, \quad c\left(q^M(\alpha)\right) = \left(\frac{1}{4\alpha}\right)^{\frac{\alpha}{\alpha-1}} \to 0$$

which imply the following limit for monopolist surplus

$$W^M = \left[\int_{\frac{1}{2}}^{1} \left(\theta q^M - \frac{1}{2} q^M\right) \mathrm{d}\theta + \int_{\frac{1}{2}}^{1} \frac{1}{2} q^M \mathrm{d}\theta\right] - c\left(q^M\right) \longrightarrow 1 \int_{\frac{1}{2}}^{1} \theta \mathrm{d}\theta - 0 = \frac{3}{8}$$

Equilibrium play under duopoly is

$$H_{2,\alpha}(q) = \mathbb{I}\left\{q \in \left[0, \left(\frac{1}{4\alpha}\right)^{\frac{1}{\alpha-1}}\right]\right\} \cdot (4\alpha q)^{\alpha-1}$$

which converges in probability to 1, so that

$$W_{2,\alpha} = \frac{1}{8}\mathbb{E}_{2,\alpha}[\boldsymbol{x} + 3\boldsymbol{y}] \to \frac{1}{8}[1+3] = \frac{1}{2}$$

Repeating the same steps we can show $W_{n,\alpha} \to \frac{1}{2}$ for each n, any intensity of competition allocates (approximately) quality 1 to all types, delivering surplus $\frac{1}{2}$, which exceeds monopoly surplus by $\frac{1}{8}$. Notice this conforms with equation (1.34), as $\frac{1}{8}$ is exactly the (limit) damaging inefficiency induced by excluding (at all quality levels) types $\left[0, \frac{1}{2}\right]$

$$\int_0^{q_\infty^M} d(q)\,\mathrm{d}q = \int_0^1 \left(\int_0^{\frac{1}{2}} \mathrm{d}\theta\right)\mathrm{d}q = \frac{1}{8}$$

1.3.4 Equilibrium properties and stability

This section has two purposes: *i*) to compare equilibrium outcomes with those obtained in models that have the most similar competition structure, and *ii*) to discuss equilibrium stability as robustness to deviations from idle firms that may unexpectedly occur as the game unfolds.

Similar models of competition

Making producers commit through an ex-ante irreversible investment is a modeling device used since Hotelling (1929): it makes equilibrium existence less problematic (and in our case also guarantees tractability) by granting incumbents some market power at the pricing stage. Separation of the production and pricing stages prevents indeed potential entrants from exploiting profitable deviations that may emerge when the competition environment or its outcome realize.

Two classical papers study similar competition games. Kreps and Scheinkman (1983) have firms commit to a quantity level before Bertrand competing (without screening) on the realized investments. The two stage equilibrium yields Cournot competition outcomes. The obvious difference with the setting of this paper is that the social value of aggregate production is obtained by summing quantities but taking the maximum over qualities of an homogenous good. As production along multiple lines is always wasteful, the result that competition may be beneficial is to some degree surprising. Champsaur and Rochet (1989) have the most similar setup as they analyze a MR duopoly where each competitor costlessly commits to a subset of qualities and then chooses a pricing function (paying the distribution costs at this stage). In committing to a quality range firms face a trade-off: they want a broad quality range to discriminate among consumers, but they also want to differentiate their products from those of the competitor as price competition lowers profit margins on neighboring qualities. They show that at a Nash equilibrium where each firm makes positive profits, the quality sets to which firms commit are always disjoint. Our investment game makes, in the language of Champsaur and Rochet (1989), firms commit to a quality range of the type $[0, q]$, so it is technologically impossible that two firms acquire disjoint sets. Indeed in all equilibria only one firm realizes positive revenues (Proposition 1.6), and first stage equilibria with active competition are only mixed (Proposition 1.7 *ii*)).

Stability

We thus far fixed an extensive form game and studied its properties: for sake of tractability we implicitly imposed strong timing and information rigidities, that must be justified by looking at, say, the length and transparency of R&D processes and patenting in the relevant markets, entry regulations etc. Contrary to Champsaur and Rochet (1989) our quality commitment stage is not cheap talk but requires a real and costly investment, though it still excludes plausible production deviations. The perfect information assumption is self-explanatory and so is the realism of its empirical counterparts.[46] We may be interested in investigating equilibria that are induced by different specifications of the competition environment. Unfortunately, no alternative specification delivered tractable results and we can only make an informal discussion of the stability of the equilibria we found. We loosely define stability as robustness to unanticipated deviations from inactive firms: the game is augmented by allowing firms that were idle to take some actions as the game unfolds. Since the game is two stage, two natural notions of stability emerge, depending on the stage at which outsiders are allowed to act.

[46] Otherwise timing is vacuous a competitors equivalently randomize over pricing functions, which would be a perfectly valid but intractable (for us) equilibrium concept.

- Interim stability: after the first stage is over a potential entrant observes the realized vector of entry qualities $\boldsymbol{q}$ and chooses whether (and eventually at which quality) to enter and play the second stage against $\boldsymbol{q}$.
- Ex-post stability: after the whole game is played a potential entrant observes the realized market pricing function and chooses whether (and eventually with which pricing function) to enter and compete with the realized contract.

Definition 1.3. The degree of interim (ex-post) stability of an equilibrium is the probability that the realized entry vector (market pricing function) does not induce interim (ex-post) entry.

Following the discussion above, it should be noticed that neither the ex-post nor the interim stability notion are associated to the equilibrium of an extended game in which active firms recognize the threat from outsiders: if he anticipates that at a later stage a potential entrant could wipe out his revenues, the interim monopolist would not (in general) choose allocation rule (1.21).[47] Similarly, the expected profit formula for the investment game would be different from (1.26) if active firms anticipated interim (and ex-post) deviations. However, if firms were playing a pure strategy equilibrium in the first stage, then the ex-post stability refinement would collapse to the Rothschild and Stiglitz (1976) equilibrium. Rothschild and Stiglitz (1976) invoke a notion of free-entry to justify the assumption that all firms observe the offered contracts[48] and must have no incentive to deviate, which is exactly what ex-post equilibrium in pure strategies requires.[49]

In this setting a Rothschild and Stiglitz (1976) contract would be a pricing function p_i. The cost of offering contract p_i is

$$\overline{c}\left(p_i\right) = c\left(\sup\left\{q : p_i\left(q\right) < \infty\right\}\right)$$

which, contrary to the original setting, does not depend on competitors' actions and consumers' demand. We associate to each set of pricing functions $\{p_i\}_{i\in\mathbb{N}}$ the firm specific revenue function $\overline{R}_j\left(\{p_i\}_{i\in\mathbb{N}}\right)$ given by (1.20) so we can define

[47]In particular, he would not alter his pricing decision if and only if this decision induces the potential deviator to abstain.

[48]In their case, a pair of coverage rate and deductible.

[49]Pure strategies make indeed the sequential nature of the game immaterial: as each player knows the realization of the quality vector and the strategy of opponents conditional on each quality vector, he knows the realization of the pricing functions and best responds to them.

Definition 1.4. An ex-post (Rothschild and Stiglitz (1976)) pure strategy equilibrium is a profile of contracts $\{p_i^\star\}_{i\in\mathbb{N}}$ such that

$$\overline{R}_i\left(p_i^\star, p_{-i}^\star\right) - \overline{c}\left(p_i^\star\right) \geq \overline{R}_i\left(p_i, p_{-i}^\star\right) - \overline{c}\left(p_i\right), \qquad \forall i, p_i$$

The following proposition establishes non-existence of ex-post equilibria and characterizes the degree of stability of monopolist and competitive equilibria.

Proposition 1.8. *i*) *There is no ex-post equilibrium in pure strategies.*

ii) *The degree of interim stability is always larger than the degree of ex-post stability. The monopolist equilibrium is interim fully stable (degree 1) but ex-post fully unstable (degree 0).*

iii) *All competitive equilibria feature intermediate degrees of interim and ex-post stability, and both of them are decreasing in* n*.*

The intuition for *i*) is the following: the monopolist pricing function (and all others abstain) is, by the same arguments used in Proposition 1.7, the unique candidate. However this time potential entrants can earn revenues that are ϵ close to $V\left(q^M\right)$ by just granting a small and equal discount to all types. The idle monopolist has no way to fight back, which breaks the candidate equilibrium.

We then notice that interim or ex-post best response implies producing a quality that lies in the doubleton set $\left\{0, q^M\right\}$. As

$$0 = R\left(q^M, q^M\right) < c\left(q^M\right) < R\left(q^M, 0\right) = V\left(q^M\right)$$

and $R\left(q^M, x\right)$ is monotonically decreasing there will be a threshold $m^\star \in \left(0, q^M\right)$ such that $R\left(m^\star, q^M\right) = c\left(q^M\right)$. We notice that at interim stage entry occurs if and only if the best active quality is below $m^\star$, while at an ex-post stage entry occurs if and only if the quality offered for free is below $m^\star$. So the degree of interim stability is the CDF (under equilibrium play) of the best quality evaluated at $m^\star$, and the degree of ex-post stability is the CDF of the second quality evaluated at $m^\star$. From this fact and the FOSD order of equilibrium statistics proved in Proposition 1.2 *i*), all results listed in *ii*) and *iii*) follow.

1.4 The Market for Information

This section discusses information markets as a potential application of the model. We highlight the implicit assumptions on the demand side and the production technology of information sources that make the framework of this paper more suited than other approaches (which are briefly reviewed) to analyze some phenomena in those markets. The main novelty is that production of information is decentralized and that the technology to convert the factor of production (attention) into state-signal structures must be taken as a primitive. We present a simple but exact microfoundation of the model in which firms observe increments of a common Brownian motion and sell the realization to agents solving a standard location problem. A natural extension in which correlation of the primary information sources induces product heterogeneity is presented since the basic model has the counterfactual implication that only one firm will be active (provide a non-trivial screening contract and make profit) even in the competitive setting. We motivate future extension of the model presented in this paper by discussing limitations imposed by the current simple structure.

1.4.1 A market for hard information

Many economically relevant situations can be modeled as a decision problem (or a game) preceded by an information acquisition phase in which agents speculate about some characteristics of the environment they will be acting in. To this aim it is necessary to specify what type of information agents can access and at what cost. The model presented in this paper can be applied to information markets with the following characteristics

- Decision makers (consumers) have no way to create their own information structure and must rely exclusively on the opinions sold by a set of profit maximizing sellers. This contrasts with models of unrestricted information acquisition with statistical pricing of information structures (e.g. Shannon entropy) which is the approach taken in (most) rational inattention models.[50]

[50] The rational inattention literature originates from the idea of Sims (1998, 2003) that decision makers are finite capacity information channels, unable to process all the information available. The information acquisition problem is equivalently rewritten by making agents pay an attention cost that is linear in the reduction of Shannon entropy, where the (per-bit) price emerges as the Lagrange multiplier on the attention constraint of the original problem. In this world, information is floating around agents that grasp it costly bit by costly bit; the fact that the "information bill" depends on some statistical property of the joint state-action distribution pays off in terms of a great tractability which accounts for some of the success of this approach.

- Firms are endowed with a technology to produce primary information structures and a (free) technology to Blackwell garble those structures. Both production and damaging occurs along a single dimension, exogenously given.

- Firms make revenues from selling information structures.[51] Price competition may become a second order concern (relative to attention and revenues from advertising) when information means entertainment and ideology drives agents' choice among opinion outlets. Newspapers, cable tv and similia are therefore not examples of information markets whose functioning is likely captured by this paper's model.[52]

- Communication of the signal (of any quality) is a costless operation that occurs without frictions: sellers have no disutility in repeating statements like "This asset is going to default with a likelihood $p \in A$." to whomever wants to pay for that, and investors have no difficulty in understanding what such statements mean,[53] to do the proper (bayesian) updating and to infer the value of such information from the decision problem they face.

- Reputation issues are also neglected: the seller cannot misreport the precision of the (menu of) signal he sells, and buyers believe the products they buy are indeed draws from the promised signal structure.[54]

Real world examples that approximately fit the description are websites that sell weather forecasts and assessments of the likelihood of default of a fixed income security (credit rating). Although (possibly) of limited interest, weather forecast is an insightful example since both the signal structure (probability of rain, temperatures

[51]The vast literature on bayesian persuasion (initiated by Kamenica and Gentzkow (2011)) and, more in general, of information design (see Bergemann and Morris (2017)) study the problem of a principal who knows the state (has already "produced" the best information structure) and optimally transmits it to a set of agents. In both cases that objective is not the maximization of revenues from selling information securities: principal's utility depends on agents' actions which he influences by tailoring the information transmission.

[52]Galperti and Trevino (2017) endogenizes the supply of information as the outcome of competition among potential information sources that choose where to locate on the accuracy/clarity space in a Myatt and Wallace (2011) setting. In a different setting Perego and Yuksel (2018) studies competitive provision and endogenous acquisition of political information with horizontal differentiation of potential consumers. In both those papers firms compete for the attention of their consumers, which is justified as many information companies make most revenues from advertising.

[53]Again, this contrasts with a rational inattention setting.

[54]Since posting a menu of prices is essentially cheap talk and the model is static, this is also a strong assumption. Reputational issues in information transmission are studied by Wang (2009) and Ottaviani and Sørensen (2006) among others.

bounds...), and the cost of acquisition (strengths of instruments, stations installed) have a clear interpretation. The website Accuweather offers basic, premium and professional subscriptions respectively for free, at $7.95 and at $19.95 monthly rate. Better packages add to basic service a longer horizon (up to 90 days), finer (hourly) weather forecast, experts opinions and radar images that clearly satisfy the feature of free damaging. The market for financial information is much larger and relevant, though some of its complexities require extensions of the model that we will address later. For the moment we let information production be the effort to evaluate the likelihood that a certain fixed-income security defaults. Credit rating has a clear signal structure (intervals of likelihoods of the default event) and to some extent can be modeled as the costly conjecturing exercise of some experts in the consultancy sector. Damaging is performed by coarsening the rating partition or hiding some parts of the report.

The key primitives that characterize such markets are *i*) the expression for the value of information and *ii*) the set of information structures that can be acquired and obtained through damaging. We proceed and give an example of both.

Information value

In an information market a type $\theta \in \Theta$ generically corresponds to a bayesian decision problem (action set, priors, utility) defined over a common uncertainty space (over which producers construct signal structures). This paper took this object as a primitive which means we can describe the shape of the optimal contract for classes of decision problems/type heterogeneity that induce a value of information having shape (1.1). It is however clear that to study specific phenomena we cannot take the value of available signal structures as exogenous but derive it from the relevant decision problem. Two recent papers that focus on information distribution derive the information value function from either heterogeneity in the decision problems or strategic externalities: Bergemann et al. (2018a) obtain a (piecewise) linear value of information when agents' types are their prior beliefs over a finite dimensional state space. They show (among many other results) that it will never be optimal for the seller to damage information by reducing precision (i.e. a quality dimension along which preferences are linear). This conforms with the "no-haggling" result stated in this paper as Corollary 1.1. In general they show that information is degraded and sold in non-trivial screening packages by revealing only a portion of the available data to the buyer: along this deterioration margin consumers' valuation are not linear.[55] Kastl et al. (2018) study the problem of a monopolist seller that may want to supply imprecise information to competitive firms that are uncertain about the marginal cost type of their contractors;[56] in their setup the state space is binary but information structures are allowed to be asymmetric and characterized by a two dimensional vector (α, β).

We can obtain an information value with concave component as in (1.1) from a decision problem that is standard in the literature: agents choose a location a to minimize the realized euclidean distance from an unknown state

$$u(a, \omega) = -(a - \omega)^2$$

so the value of an information structure $\mathcal{S} : \Omega \to \Delta(S)$ is just the expected (i.e. after observing a draw from $\mathcal{S}$) reduction in the variance

$$g(\mathcal{S}) = -\mathbb{E}_{\mathcal{S}}[\mathbb{V}_{post}] + \mathbb{V}_{prior}$$

If agents have normal prior belief $\omega \sim \mathcal{N}\left(\mu, \tau_p^{-1}\right)$ and information structures take the form $\mathcal{N}(\omega, q^{-1})$ for $q \in \mathbb{R}$,[57] then such value is given by

$$g(q) = -\frac{1}{q + \tau_p} + \frac{1}{\tau_p} = \frac{q}{\tau_p(q + \tau_p)} \tag{1.38}$$

which is a concave function in q. Moreover, notice that

$$g''(q) = -\frac{1}{(q + \tau_p)^2}$$

is decreasing (in absolute value) in τ_p, so larger precision implies "more concave" returns from quality in the sense of footnote 15. To get the specification (1.1) we should therefore assume that agents solve two independent "location problems" and have heterogeneous valuations from guessing right the problem with larger prior variance. In general, we can keep the common location game as inducing the concave

[55] As a concrete example of damaging through partial revelation occurring in information markets, Bergemann et al. (2018b) point at the "Undisclosed Debt monitoring" packages sold by Equifax in which the data broker offers individual rating reports to financial firms considering application for loans in three different versions differing in the number of "red flags" that the lender receives if the borrowers' history includes some negative events.

[56] They focus on the trade-off of a monopolist that may want to sell imprecise signal in order to limit distortions due to internal agency problems.

[57] Induced by speculation effort (1.39) (next section).

component, and take the linear part as some Taylor approximation for an additional type-dependent returns from precision.

Information production

Primary information structures are the state-signal correlations that can be produced by the firms. Suppose $\Omega = \mathbb{R}$ and the choice of primary information corresponds to observing a Brownian motion

$$dX_t = \omega dt + dW_t^j \tag{1.39}$$

for some period of (costly) time. After staring at the Brownian Motion for q_j units of time, firm j "produced" a signal (sufficient for ω)

$$s_j = \frac{1}{\sqrt{q_j}} X_{q_j} \sim \mathcal{N}\left(\omega, \frac{1}{q_j}\right)$$

about the state, which can be (damaged and) distributed to interested parties. Clearly, as every firm $i \in \mathbb{N}$ observes the same Brownian motion but just for different time, whoever stares at it the longer can push out of the market (owns a product that is superior to) everyone else. Every firm j can "quote" a market for any precision that is below q_j since inferior qualities can be obtained by reporting the Brownian motion at a period before q_j or by adding independent normal noise.[58] Notice we have made the implicit assumption that damaged structures must fall in the normal family. However there is no reason for which sellers should have this restriction: they can report whether it lies in a certain region, whether it is closer to point A or B, or commit to any garbling. The restriction to single dimensional quality is therefore substantial especially when it comes to damaging, and becomes untenable when we have a generic state space without parametric restrictions on primary structures.[59]

[58]If agents do not care about correlation (e.g. they play independent decision problems) then how the signal is damaged is immaterial. If they were playing a game then also the correlation of the signal would induce a value. In particular, in presence of strategic substitutability agent derive value from being uncorrelated and we would obtain the result that a monopolist may acquire a quality higher than the highest quality distributed, just to be able to damage it in an agent-independent way.

[59]It is somehow natural to have single dimensional production sets: we can say that a weather forecaster can buy a stronger telescope that allows for fixed maps into signal structures, or that financial experts can only determine the intensity of their search and the resulting Markov kernels are model primitives. It is clearly a much stronger assumption to say that such conjecturing effort can be damaged only along the production dimension.

Correlation as product heterogeneity

Proposition 1.6 gives the implication that only one firm sells positive qualities and makes profits. This seems to be counterfactual at least in the market for financial consultancy where many people make money out of saying something. Product homogeneity in this framework arises from having all firms observe the same Brownian Motion (1.39): the empirical counterpart of this assumption is that all experts look at the same set of evidence (share a common reasoning process), or that meteorological instruments make perfectly correlated errors.

Such extreme assumption can be relaxed by letting the Wiener process driving observation of firm j be

$$dW_t^j = \rho dW_t + \sqrt{1-\rho^2} dZ_t^j \tag{1.40}$$

where dZ_t^j is a firm-specific process[60] and $\rho \in [0,1]$ parametrizes the correlation of the conjecturing effort of the different firms. Suppose for expositional clarity that there are two firms i, j and that at the investment stage they produced $q_i > q_j$. Despite being of inferior quality (correlation with the state) now X_{q_j} contains information about ω even after conditioning on X_{q_i}. Indeed, it holds

$$Cov\left(X_{q_j}, X_{q_i}\right) = \rho \min\left\{q_i, q_j\right\} = \rho q_j$$

Consumers only care about the final precision of the signal they observe (see (1.38)), that is $u\left(q_i, q_j, \theta\right)$ admits the aggregator representation

$$u\left(q_i, q_j, \theta\right) = u\left(\Psi\left(q_i, q_j\right), \theta\right)$$

where $\Psi : Q^2 \to Q$, derived by the normal updating formulas is the piecewise convex function given by

$$\Psi\left(q_i, q_j\right) = \frac{q_i\left(q_j\left(1-2\rho\right)+q_i\right)}{q_i - q_j\rho^2}$$

Notice that if $\rho = 0$ (signals are uncorrelated), then $\Psi(q) = q_i + q_j$ and we have a model of additive social value of production as in Kreps and Scheinkman (1983), though the distribution problem is subject to the screening frictions. If $\rho = 1$ then we are back to homogenous products and maximum aggregator $\Psi(q) = q_i$ which has been the subject of study of this paper. Continuity of Ψ implies full deterioration domain,

[60] Meaning dW_t, dZ_t^j, dZ_t^i are uncorrelated Wiener processes.

namely that that for each $\{q_i\}_{i\in\mathbb{N}}$ and $q' < \Psi\left(\{q_i\}_{i\in\mathbb{N}}\right)$ we can find $\{q_i'\}_{i\in\mathbb{N}} \leq \{q_i\}_{i\in\mathbb{N}}$ so that $q' = \Psi\left(\{q_i'\}_{i\in\mathbb{N}}\right)$.

We define the cost of aggregate quality as the value function

$$\bar{c}(q) = \min_{\{q_i\}_{i\in\mathbb{N}}} \sum_{i\in\mathbb{N}} c(q_i), \qquad \text{s.t. } \Psi\left(\{q_i\}_{i\in\mathbb{N}}\right) \geq q \tag{1.41}$$

With the full deterioration property we can adapt the techniques used for the homogenous product and characterize first and second best. Second best is equivalent to a monopolist that owns all production sources and distributes damaged qualities only subject to information frictions: he allocates packages $\{q_i(\theta)\}_{i\in\mathbb{N},\theta\in\Theta}$ subject to IR and IC where each type θ can only choose among profiles $\{q_i(\theta')\}_{i\in\mathbb{N}}$ (cannot pick $q_i(\theta')$ and $q_j(\theta'')$ for $i \neq j$). Proposition (1.1) applies verbatim to characterize the second best distribution of aggregate qualities which are produced at cost (1.41). We cannot however decentralize the second best as a pricing stage equilibrium among competitive firms. The technical complications to solve for a competitive equilibrium with screening and heterogeneous products are illustrated in the Handbook chapter of Stole (2007), and we could not extend the tractability of competitive equilibria for homogenous digital goods to this more general case.

1.4.2 Limitations and future research

Section 1.4 had three objectives: *i*) discuss features of information markets that make them fit to be studied under the framework of this paper rather than under alternative approaches, *ii*) use standard building blocks to provide a microfoundation for an information market that fits exactly the reduced form description given in the paper, and *iii*) suggest correlation in primary structures as a natural way to introduce product heterogeneity in a market that fails one stark empirical implication of the homogenous product model. Section explored a tractable way to allow for heterogeneous in preferences that admit a quality aggregator. The unsolved technical challenge is the specification of a competition environment that delivers tractable equilibria. We conclude the section exploring other extensions of the model

- Contrary to a growing literature on information acquisition in markets, we do not allow for strategic interaction at the decision stage as this would make an exogenous (i.e. independent of equilibrium play) specification of the value for information untenable.[61] The extension would make a second quality dimension (broadly speaking, correlation with information given to co-players) emerge endogenously. Broadly speaking, this direction points at an heterogeneous agent version of Myatt and Wallace (2011) (or derived setups) with information acquisition as in (1.40) and price competition.

- The quality restrictions become substantive when the state space is large: as investors have interests that are differentiated either geographically or for the type of assets they trade, restricting sellers' marketable products to be single dimensional is unreasonable. We could think of a model with horizontally differentiated agents that care, say, only about some dimension of a large state space. A conjecturing effort produces signal structures about the whole state space but firms may choose to sell what they know about different portions of the state space as different products. We could use this specification to ask under what conditions firms favor production and marketing of signal structures characterized by a large breadth (learn about the approximate location of many states) or depth (focus on one state and identify it more precisely) component.

- The issue of non-excludability is particularly relevant for information markets: beyond prohibiting re-selling of opinions, private information may be "leaked" through aggregate variables (this channel is explored in Admati and Pfleiderer (1986)). Many financial information packages include agent specific information (as the rating of potential borrowers in the Equifax example of Bergemann et al. (2018b)), and live prices (Bloomberg vs Reuters), which somehow reduce the concern of failure of non-excludability.

1.5 Conclusions

In this paper we developed a model of production and distribution of digital goods. The monopolist problem reduces to quality screening where the cost of an

[61]Several contributions (among which Hellwig and Veldkamp (2009), Myatt and Wallace (2011), and Colombo et al. (2014)) show that in the presence of strategic externalities also the information acquisition game has strategic complementarities (i.e. there is an information acquisition externality). The game structure makes the value of information an endogenous object, but the cost is still exogenous even in cases where the supply side is a rich set of information sources who are characterized by an accuracy-clarity pair (as in Myatt and Wallace (2011)).

allocation is not additively separable but depends solely on one statistic of such allocation (the maximum). Under regularity assumptions on the demand side the optimal allocation is characterized by a bunching at the top threshold that increases with the quality cap, which is then easy to solve for. Market power and asymmetric information induce interdependent inefficiencies in acquisition and distribution of the digital good. Preventing damaging always worsens the acquisition inefficiency and induces exclusion of types that would be served under the unconstrained contract. The mechanical undoing of damaging inefficiencies counters those perverse effects yielding an ambiguous welfare impact of the policy whose sign depends on the cost primitive. We then studied competition in digital goods markets as an extensive form game in which investment in quality is a sunk cost at the pricing stage. Monopoly is the only equilibrium in pure strategies, but there are also equilibria with different levels of competition. Competition induces wasteful double spending and worsens underacquisition since the highest quality distributed by a competitive market is below monopolist level. However, Bertrand forces induce a contraction of the screening domain that reduces distributional inefficiencies. Across equilibria with active competition the duopoly is Pareto dominant and welfare is decreasing in the intensity of competition. The welfare comparison between monopoly and duopoly is ambiguous and we can tailor the cost function to completely shut down the channels that favor either of them. The monopolist equilibrium is always subject to ex-post deviations while the duopoly features the highest degree of ex-post stability across all equilibria. We concluded by discussing how to apply and extend the model to study some phenomena in the market for information from a novel perspective.

References

Anat R Admati and Paul Pfleiderer. A monopolistic market for information. *Journal of Economic Theory*, 39(2):400–438, 1986.

Paul Beaudry and Michel Poitevin. Competitive screening in financial markets when borrowers can recontract. *The Review of Economic Studies*, 62(3):401–423, 1995.

Dirk Bergemann and Stephen Morris. Information design: A unified perspective. 2017.

Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–57, 2015.

Dirk Bergemann, Alessandro Bonatti, and Alex Smolin. The design and price of information. *American Economic Review*, 108(1):1–48, 2018a.

Dirk Bergemann, Alessandro Bonatti, et al. Markets for information: An introduction. *Cowles Foundation for Research in Economics Discussion paper*, 2018b.

Alberto Bisin and Piero Gottardi. Competitive equilibria with asymmetric information. *Journal of Economic Theory*, 87(1):1–48, 1999.

Alberto Bisin, John D Geanakoplos, Piero Gottardi, Enrico Minelli, and Herakles Polemarchakis. Markets and contracts. *Journal of Mathematical Economics*, 47(3): 279–288, 2011.

Paul Champsaur and Jean-Charles Rochet. Multiproduct duopolists. *Econometrica: Journal of the Econometric Society*, pages 533–557, 1989.

Luca Colombo, Gianluca Femminis, and Alessandro Pavan. Information acquisition and welfare. *The Review of Economic Studies*, 81(4):1438–1483, 2014.

Raymond J Deneckere and R Preston McAfee. Damaged goods. *Journal of Economics & Management Strategy*, 5(2):149–174, 1996.

Pradeep Dubey and John Geanakoplos. Competitive pooling: Rothschild-stiglitz reconsidered. *The Quarterly Journal of Economics*, 117(4):1529–1570, 2002.

Bhaskar Dutta and Rajiv Vohra. Incomplete information, credibility and the core. *Mathematical Social Sciences*, 50(2):148–165, 2005.

Simone Galperti and Isabel Trevino. Coordination motives and competition for attention in information markets. Technical report, Mimeo, 2017.

Daniel Garrett, Renato Gomes, and Lucas Maestri. Competitive screening under heterogeneous information. 2014.

Andrew V Goldberg and Jason D Hartline. Envy-free auctions for digital goods. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 29–35. ACM, 2003.

Andrew V Goldberg, Jason D Hartline, and Andrew Wright. Competitive auctions and digital goods. In *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 735–744. Society for Industrial and Applied Mathematics, 2001.

Andrew V Goldberg, Jason D Hartline, and Andrew K Wright. Systems and methods for pricing and selling digital goods, January 10 2006. US Patent 6,985,885.

Jong-Hee Hahn. Damaged durable goods. *The RAND Journal of Economics*, 37(1): 121–133, 2006.

Oliver D Hart and Jean Tirole. Contract renegotiation and coasian dynamics. *The Review of Economic Studies*, 55(4):509–540, 1988.

Jason D Hartline and Tim Roughgarden. Optimal mechanism design and money burning. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 75–84. ACM, 2008.

Christian Hellwig and Laura Veldkamp. Knowing what others know: Coordination motives in information acquisition. *The Review of Economic Studies*, 76(1):223–251, 2009.

Martin F Hellwig. A note on the specification of interfirm communication in insurance markets with adverse selection. *Journal of Economic Theory*, 46(1):154–163, 1988.

H. Hotelling. Stability in competition. *Economic Journal, 39*, 1929.

Roman Inderst. Durable goods with quality differentiation. *Economics Letters*, 100 (2):173–177, 2008.

Gerald D Jaynes. Competitive screening and market segmentation. 2006.

Gerald David Jaynes. Equilibria in monopolistically competitive insurance markets. *Journal of Economic Theory*, 19(2):394–422, 1978.

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

Jakub Kastl, Marco Pagnozzi, and Salvatore Piccolo. Selling information to competitive firms. *The RAND Journal of Economics*, 49(1):254–282, 2018.

David M Kreps and Jose A Scheinkman. Quantity precommitment and bertrand competition yield cournot outcomes. *The Bell Journal of Economics*, pages 326–337, 1983.

Eric Maskin and John Riley. Monopoly with incomplete information. *The RAND Journal of Economics*, 15(2):171–196, 1984.

R Preston McAfee. Pricing damaged goods. 2007.

Michael Mussa and Sherwin Rosen. Monopoly and product quality. *Journal of Economic theory*, 18(2):301–317, 1978.

Shigeo Muto. An information good market with symmetric externalities. *Econometrica: Journal of the Econometric Society*, pages 295–312, 1986.

David P Myatt and Chris Wallace. Endogenous information acquisition in coordination games. *The Review of Economic Studies*, 79(1):340–374, 2011.

Nick Netzer and Florian Scheuer. Competitive screening in insurance markets with endogenous wealth heterogeneity. *Economic Theory*, 44(2):187–211, 2010.

Marco Ottaviani and Peter Norman Sørensen. Reputational cheap talk. *The Rand journal of economics*, 37(1):155–175, 2006.

Martin Peitz and Patrick Waelbroeck. Piracy of digital products: A critical review of the theoretical literature. *Information Economics and Policy*, 18(4):449–476, 2006.

Jacopo Perego and Sevgi Yuksel. Media competition and social disagreement. 2018.

Arnold Polanski. A decentralized model of information pricing in networks. *Journal of Economic Theory*, 136(1):497–512, 2007.

Jean-Charles Rochet and Philippe Choné. Ironing, sweeping, and multidimensional screening. *Econometrica*, pages 783–826, 1998.

Michael Rothschild and Joseph Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The quarterly journal of economics*, pages 629–649, 1976.

Christopher A Sims. Stickiness. In *Carnegie-Rochester Conference Series on Public Policy*, volume 49, pages 317–356. Elsevier, 1998.

Christopher A Sims. Implications of rational inattention. *Journal of monetary Economics*, 50(3):665–690, 2003.

Michael Smart. Competitive insurance markets with two unobservables. *International Economic Review*, 41(1):153–170, 2000.

Padmanabhan Srinagesh and Ralph M Bradburd. Quality distortion by a discriminating monopolist. *The American Economic Review*, pages 96–105, 1989.

Lars A Stole. Price discrimination and competition. *Handbook of industrial organization*, 3:2221–2299, 2007.

Lisa N Takeyama. The welfare implications of unauthorized reproduction of intellectual property in the presence of demand network externalities. *The journal of industrial economics*, pages 155–166, 1994.

Hal R Varian. Buying, sharing and renting information goods. *The Journal of Industrial Economics*, 48(4):473–488, 2000.

Rajiv Vohra. Incomplete information, incentive compatibility, and the core. *Journal of Economic Theory*, 86(1):123–147, 1999.

Hefei Wang. Reputation acquisition of underwriter analysts–theory and evidence. *Journal of Applied Economics*, 12(2):331–363, 2009.

Robert Wilson. Information, efficiency, and the core of an economy. *Econometrica: Journal of the Econometric Society*, pages 807–816, 1978.

Robert B Wilson. *Nonlinear pricing*. Oxford University Press on Demand, 1993.

Proof Appendix

1.5.1 Proof of Proposition 1.1 and Remark 1.1

As $u_q > 0$ for any type θ, deterioration is inefficient: the planner produces a quality that equates the average marginal valuation to the marginal cost and distributes such quality to each type. (1.5) is the first order condition of problem (1.4) after noticing the efficient allocation is flat, whose sufficiency is immediate.

A seller with perfect information can charge a type-dependent price p_θ. Then the profit maximization problem coincides with the social surplus problem, he produces $q^\star$, distributes it to all types and he extracts all the surplus.

1.5.2 Proof of Lemma 1.1

Consider the monopolist problem

$$\begin{array}{rcc} \max_{\rho,p:\Theta\to Q\times\mathbb{R}} & & \int_\Theta p(\theta)\,\mathrm{d}F(\theta) - \overline{c}(\rho) \\ & \text{s.t.} & \text{IC,IR} \end{array} \tag{1.42}$$

and define

$$\omega(x) = \{\rho, p : \Theta \to Q\times\mathbb{R} : \text{ IC, IR hold and } \overline{c}(\rho) \le x\}$$

be the set incentive compatible and individually rational allocation and pricing function whose cost does not exceed x. Using this constraint sets, problem (1.42) can be rewritten as

$$\begin{aligned} &\max_{x,\{p:\exists\rho,(p,\rho)\in\omega(x)\}} \int_\Theta p(\theta)\,\mathrm{d}F(\theta) - x \\ &= \max_{x\in\mathbb{R}} \left[\max_{\{p:\exists\rho,(p,\rho)\in\omega(x)\}} \int_\Theta p(\theta)\,\mathrm{d}F(\theta)\right] - x \end{aligned} \tag{1.43}$$

now given the specification of the cost function (1.6)

$$\begin{aligned} \rho\in\omega(x) \quad \text{only if} &\quad c\left(\max_\theta \rho(\theta)\right) \le x \\ \iff &\quad \max_\theta \rho(\theta) \le c^{-1}(x) \end{aligned}$$

so

$$\omega(x) = \left\{\rho, p : \Theta \to Q \times \mathbb{R} : \text{ IC, IR hold and } \max_{\theta} \rho(\theta) \leq c^{-1}(x)\right\}$$

and therefore

$$\max_{\{p : \exists \rho, (p,\rho) \in \omega(x)\}} \int_{\Theta} p(\theta)\, \mathrm{d}F(\theta) = V\left(c^{-1}(x)\right)$$

where V is the value of quality defined in (1.8) as the constraint set of that problem coincides with $\omega\left(c^{-1}(x)\right)$. Redefining the domain of choice to be $Q = c^{-1}(\mathbb{R})$ (which we can do as c strictly increasing), problem (1.43) becomes

$$\max_{q \in Q} V(q) - c(q)$$

as we wanted to show.

1.5.3 Proof Proposition 1.2

Fix an arbitrary q. The monopolist chooses allocations and transfers $\{q_i, T_i\}_{i \in \{H,L\}}$ to solve

$$V(q) = \max_{\{q_i, T_i\}_{i \in \{H,L\}}} (1 - \pi) T_1 + \pi T_2$$

subject to

$$u(q_i, i) - T_i \geq u(q_j, i) - T_j$$

$$u(q_i, i) - T_i \geq 0$$

$$q_i \leq q$$

Combining the incentive constraint for high and low types we get monotonicity of the optimal allocation. Optimality implies the rationality constraint of the low type and the incentive to deviate from high to low types must be binding. With those observations the revenue maximization problem can written as a control problem where we choose allocation to high type and low type, x, y respectively.

$$V(q) = \max_{0 \leq y \leq x \leq q} u(L, y) + \pi\left(u(H, x) - u(H, y)\right)$$

Maximization with respect to x immediately gives the corner solution

$$x^{\star} = q$$

Necessary condition for interior y is

$$u'(L, y^{\star}) - \pi u'(H, y^{\star}) = 0$$

$$\frac{g'(q) + \theta_L}{g'(q) + \theta_H} = \pi \tag{1.44}$$

Notice that by assumption $u''_H(q) = g''(q) = u''_L$ so

$$\frac{u'(L, x)}{u'(H, x)} = \frac{g'(x) + \theta_L}{g'(x) + \theta_H}$$

is monotonically decreasing, and equation (1.44) has (at most) one solution. and also guarantees that the SOC

$$u''(L, y^{\star}) - \pi u''(H, y^{\star}) = (1 - \pi) g''(y^{\star}) < 0$$

is satisfied at the critical point, giving the unique solution to the program.

Now, if $y^{\star} \le q$, then also the monotonicity constraint is satisfied and we have a global solution. If the threshold $y^{\star}$ is below the maximal quality we need to compare the fully pooling and exclusion equilibrium. By the single crossing property, for all $x < y^{\star}$ it holds $u'(L, x) > \pi u'(H, x)$, therefore

$$u(q, L) = \int_0^q u'(L, x)\,\mathrm{d}x > \int_0^q \pi u'(H, x)\,\mathrm{d}x = \pi u(q, H)$$

and it is more profitable to serve q to all consumers at price $u(q, L)$. This proves L receives $y^{\star}$ if $y^{\star} \le q$ and q otherwise. Given allocation we can infer transfers from the binding constraints and obtain the expression for the revenue (and marginal revenue) function.

1.5.4 Proof of Lemma 1.2

We check that the virtual valuation is indeed a supermodular function in q, θ.

$$\begin{aligned} \tfrac{\partial}{\partial q \partial \theta} vv(q, \theta) &= \tfrac{\partial}{\partial q \partial \theta} \left[g(q) + \theta q - q \tfrac{1 - F(\theta)}{f(\theta)} \right] \\ &= \tfrac{\partial}{\partial \theta} \quad g'(q) + \theta - \tfrac{1 - F(\theta)}{f(\theta)} \\ &= 1 - \underbrace{h'(\theta)}_{<0} \quad \ge 0 \end{aligned}$$

Now notice that

$$g(q) + q \left[\theta - \frac{1 - F(\theta)}{f(\theta)} \right]$$

is maximized at $q = \infty$ whenever the multiplier on the linear part is greater than zero, since this is a monotonically increasing function of q. On the contrary, when $\theta \leq \widetilde{\theta}$, then $\theta - \frac{1-F(\theta)}{f(\theta)} < 0$ and the objective is concave as the difference between a concave function g and a linear part. So maximizer $\beta(\theta)$ is characterized by the first order condition

$$g'(\beta(\theta)) = \frac{1 - F(\theta)}{f(\theta)} - \theta$$

As g' is a strictly increasing function and can invert it to get the other branch of the maximizer function

$$\beta(\theta) = \begin{cases} (g')^{-1}(h(\theta) - \theta) & \theta < \widetilde{\theta} \\ \infty & \theta \geq \widetilde{\theta} \end{cases}$$

as we wanted to show.

1.5.5 Proof of Proposition 1.1

Fix a generic quality cap q. We want to show $i)$, that is $\rho(q, \theta) = \min\{\beta(\theta), q\}$.

The quality constraint $\rho(\theta) \leq q$ defining problem (1.8) is inserted in the objective function by subtracting to type dependent revenues the cost

$$c_\infty(q') = \begin{cases} 0 & q' \leq q \\ \infty & else \end{cases}$$

So the cap conditional problem equivalently reads

$$\begin{array}{ccc} V(q) & \longmapsto & \max_{\rho,p:\Theta\to Q\times\mathbb{R}} \int_\Theta p(\theta) - c_\infty(\rho(\theta))\,\mathrm{d}F(\theta) \\ & \text{s.t.} & \text{IC , IR} \end{array}$$

Notice $c_\infty(q')$ is not differentiable, but can be approximated by the continuously differentiable convex function

$$c_n(q') = \left(\frac{q'}{q}\right)^n$$

We can now define the sequence of auxiliary problems

$$\begin{array}{ccc} V_n(q) & \longmapsto & \max_{\rho,p:\Theta\to Q\times\mathbb{R}} \int_\Theta p(\theta) - c_n(\rho(\theta))\,\mathrm{d}F(\theta) \\ & \text{s.t.} & \text{IC , IR} \end{array}$$

as $\lim_{n\to\infty} c_n(q') = c_\infty(q')$, the objective in V_n converges to the objective in V and as policies and values of the auxiliary problems are bounded, the sequence of solutions (ρ_n, p_n) converges to the solution to the original problem. The auxiliary problem for a generic n is a monopolist screening problem with additively separable cost function, which we solve using standard arguments.

Firstly, the pairwise comparison of incentive constraints implies that the allocation ρ_n is monotonically increasing. Then, using the envelope theorem assuming sufficiency of the first order approach we obtain

$$\frac{\mathrm{d}}{\mathrm{d}\theta}\left[u\left(\rho_n\left(\theta\right),\theta\right)-p_n\left(\theta\right)\right]=u_\theta\left(\rho_n\left(\theta'\right),\theta'\right)=\rho_n\left(\theta'\right)$$

from which we get the payoff equivalence function

$$u\left(\rho_n\left(\theta\right),\theta\right)-p_n\left(\theta\right)=\int_0^\theta \rho_n\left(\theta'\right)\mathrm{d}\theta'$$

from which we infer prices

$$p_n\left(\theta\right)=u\left(\rho_n\left(\theta\right),\theta\right)-\int_0^\theta \rho_n\left(\theta'\right)\mathrm{d}\theta'$$

and substitute in the objective to get the relaxed problem

$$W_n\left(q\right)\longrightarrow \max_{\rho_n(\theta)\text{ increasing}}\int_\Theta\left[u\left(\rho_n\left(\theta\right),\theta\right)-c_n\left(\rho_n\left(\theta\right)\right)-\int_0^\theta \rho_n\left(\theta'\right)\mathrm{d}\theta'\right]f\left(\theta\right)\mathrm{d}\theta$$

integrating by parts we obtain

$$\max_{\rho_n(\theta)\text{ increasing}}\int_\Theta\left[u\left(\rho_n\left(\theta\right),\theta\right)-c_n\left(\rho_n\left(\theta\right)\right)-\rho_n\left(\theta\right)\frac{1-F\left(\theta\right)}{f\left(\theta\right)}\right]f\left(\theta\right)\mathrm{d}\theta$$

pointwise maximization of the integrand gives that a candidate $\rho_n\left(\theta\right)$ must satisfy

$$g'\left(\rho_n\left(\theta\right)\right)+\rho_n\left(\theta\right)-\frac{1}{h\left(\theta\right)}-\frac{n}{q^n}\left(\rho_n\left(\theta\right)\right)^{n-1}=0$$

that $\rho_n\left(\theta\right)$ so defined satisfies monotonicity follows from supermodularity of the profit function

$$r\left(x,\theta\right)=g\left(x\right)+\theta x-x\frac{1-F\left(\theta\right)}{f\left(\theta\right)}-\left(\frac{x}{q}\right)^n$$

as subtraction of a (convex) function of x does not change the sign of the cross derivative computed in Lemma 1.2 (neither sufficiency of the first order condition).

Recall that $\beta(\theta)$ solves

$$g'(\beta(\theta)) + \beta(\theta) - \frac{1}{h(\theta)} = 0$$

and as

$$\lim_{n\longrightarrow\infty} \frac{n}{q^n}(x)^{n-1} = \begin{cases} 0 & x < q \\ 1 & x = q \\ \infty & x > q \end{cases}$$

we conclude that the pointwise limit of $\rho_n(\theta)$ is

$$\rho_n(\theta) \rightarrow \begin{cases} \beta(\theta) & \text{if } \beta(\theta) < q \\ q & \beta(\theta) > q \end{cases}$$

and so $\rho(q,\theta) = \min\{\beta(\theta), q\}$, which is our desideratum.

ii) Now assume that types are uniformly distributed.[62] Using that, we write

$$V(q) = \int_0^1 u(\rho(\theta,q),\theta) - (1-\theta)\rho(\theta,q)\,\mathrm{d}\theta \tag{1.45}$$

it is convenient to define $b : Q \rightarrow \left[0, \frac{1}{2}\right]$ the inverse β function, namely

$$b(q) = \beta^{-1}(q) = \max\left\{0, \frac{1-g'(q)}{2}\right\} \tag{1.46}$$

so that

$$\rho(q,\theta) = \begin{cases} \beta(\theta) & \theta \leq b(q) \\ q & else \end{cases}$$

which can be substituted in 1.45 to obtain

$$V(q) = \int_0^{b(q)} u(\beta(\theta),\theta) - (1-\theta)\beta(\theta)\,\mathrm{d}\theta + \int_{b(q)}^1 u(q,\theta) - (1-\theta)\,q\mathrm{d}\theta$$

[62] See discussion in Section

now we can differentiate it

$$\begin{aligned} V'(q) &= b'(q)\left[u(q,b(q)) - (1-b(q))\,q\right] \\ &- b'(q)\left[u(q,b(q)) - (1-b(q))\,q\right] \\ &+ \int_{b(q)}^{1} u_1(q,\theta) - (1-\theta)\,\mathrm{d}\theta \\ &= \int_{b(q)}^{1} g'(q) + (2\theta - 1)\,\mathrm{d}\theta \\ &= (1-b(q))\,(g'(q) - 1) + \left(1 - b(q)^2\right) \\ &= (1-b(q))\left[g'(q) - 1 + (1+b(q))\right] \\ &= (1-b(q))\left[g'(q) + b(q)\right] \end{aligned}$$

substituting expression (1.46) for b we get:

If $1 < g'(q)$ then $b(q) = 0$ and $V'(q) = g'(q)$

Otherwise,

$$V'(q) = \left(\frac{1 + g'(q)}{2}\right)^2$$

which is the expression in the Proposition. We are left now left to show that V' is $\mathcal{C}^1$. Continuous differentiability in the two branches is immediate, we need to show they are smoothly pasted.

$$\left(\frac{1+g'(q)}{2}\right)^2 \Bigg|_{g'(q)=1} = 1 = g'(q)\Big|_{g'(q)=1}$$

proves continuity, while

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}q}\left(\tfrac{1+g'(q)}{2}\right)^2 \Bigg|_{g'(q)=1} &= 2\frac{1+g'(q)}{2}\frac{g''(q)}{2}\Bigg|_{g'(q)=1} \\ &= 2\frac{1+1}{2}\frac{g''(q)}{2} \\ &= g''(q) \\ &= \frac{\mathrm{d}}{\mathrm{d}q} g'(q) \end{aligned}$$

proves continuous differentiability.

iii) Per Proposition 1.1 the efficient quality $q^\star$ is determined by the first order condition

$$g'(q) + \frac{1}{2} = c'(q)$$

Monopolist quality q^M instead solves

$$V'(q) = c'(q)$$

so it is sufficient to show

$$V'(q) < g'(q) + \frac{1}{2} \tag{1.47}$$

always. Clearly, $g'(q) + \frac{1}{2} > g'(q)$ so in the full bunching region this is true. It is also easy to check that

$$x \leq 1 \implies x + \frac{1}{2} > \left(\frac{1+x}{2}\right)^2$$

proving $V'(q)$ is strictly below efficient marginal surplus even when $g'\left(q^M\right) \leq 1$. As (1.47) always holds, $q^M < q^\star$ and we have inefficient acquisition.

1.5.6 Proof of Proposition 1.3

We firstly need to compute consumers' surplus. As u_θ is $\rho(\theta, q)$, type θ welfare when the cap is q reads

$$W(\theta, q) = u(0) + \int_0^\theta u_\theta(\rho(\theta', q), \theta')\,\mathrm{d}\theta' = \int_0^\theta \min\{\beta(\theta'), q\}\,\mathrm{d}\theta' \tag{1.48}$$

Integrating over $\Theta = [0,1]$ to get total consumer surplus, we have

$$\begin{aligned} W(q) &= \int_0^1 W(\theta, q)\,\mathrm{d}\theta \\ &= \int_0^1 \int_0^\theta \min\{\beta(\theta'), q\}\,\mathrm{d}\theta'\mathrm{d}\theta \\ &= \int_0^1 \min\{\beta(\theta), q\}(1-\theta)\,\mathrm{d}\theta \\ &= \int_0^1 \min\left\{(g')^{-1}(2\theta-1), q\right\}(1-\theta)\,\mathrm{d}\theta \\ &= \int_0^{b(q)} (g')^{-1}(2\theta-1)(1-\theta)\,\mathrm{d}\theta + \int_{b(q)}^1 q(1-\theta)\,\mathrm{d}\theta \end{aligned}$$

First line is definition, second substitutes (1.7), third is integration by parts, fourth expresses $\beta(\theta)$ and finally breaks the integral in the parts above and below q.

When differentiating, as for the marginal revenues V', the terms in b' will drop so we are left with

$$\begin{aligned} W'(q) &= \int_{b(q)}^1 \frac{\partial}{\partial q} q(1-\theta)\,\mathrm{d}\theta \\ &= \int_{b(q)}^1 (1-\theta)\,\mathrm{d}\theta \\ &= (1-b(q)) - \frac{1}{2}x^2\Big|_{b(q)}^1 \\ &= (1-b(q)) - \frac{1}{2}\left(1-b(q)^2\right) \\ &= \frac{1}{2}(1-b(q))^2 \\ &= \frac{1}{2}\left(\frac{1+g'(q)}{2}\right)^2 = \frac{1}{2}V'(q) \end{aligned} \tag{1.49}$$

Integrating marginal surplus below $(\frac{1}{2})$, and above (1.49) $(g')^{-1}(1)$ we obtain (1.14).

To obtain the inefficiencies decomposition, notice first best surplus is given by

$$\begin{aligned} S^{FB} &= \tfrac{1}{2}q^\star + g(q^\star) - c(q^\star) \\ &= \int_0^{q^\star} \left(\tfrac{1}{2} + g'(q) - c'(q)\right) \mathrm{d}q \end{aligned}$$

Summing marginal revenues (1.12) and consumer surplus (1.49) we get monopolist surplus below $(g')^{-1}(1)$ grows as first best (has no damaging), while above it grows with slope

$$\left(\frac{1+g'(q)}{2}\right)^2 - c'(q) + \frac{1}{2}\left(\frac{1+g'(q)}{2}\right)^2 = \frac{3}{2}\left(\frac{1+g'(q)}{2}\right)^2 - c'(q)$$

Monopolist welfare can therefore be written as

$$W_M = \int_0^{(g')^{-1}(1)} \left(\tfrac{1}{2} + g'(q) - c'(q)\right) \mathrm{d}q + \int_{(g')^{-1}(1)}^{q^M} \tfrac{3}{2}\left(\tfrac{1+g'(q)}{2}\right)^2 - c'(q)\, \mathrm{d}q \qquad (1.50)$$

Notice in the region $\left[(g')^{-1}(1), q^M\right]$ we have marginal damaging inefficiencies

$$\begin{aligned} d(q) &= \underbrace{\frac{1}{2} + g'(q) - c'(q)}_{(S^{FB})'} - \underbrace{\left[\frac{3}{2}\left(\frac{1+g'(q)}{2}\right)^2 - c'(q)\right]}_{W_M'} \\ &= \tfrac{1}{8}\left(1 + (2 - 3g'(q))\, g'(q)\right) \end{aligned}$$

By splitting integrals in the three regions $\left[0, (g')^{-1}(1)\right]$, $\left[(g')^{-1}(1), q^M\right]$ and $\left[q^M, q^\star\right]$ we can write losses relative to first best as

$$\begin{aligned} S^{FB} - W_M = &\int_0^{(g')^{-1}(1)} \left(\tfrac{1}{2} + g'(q) - c'(q)\right) \mathrm{d}q - \int_0^{(g')^{-1}(1)} \left(\tfrac{1}{2} + g'(q) - c'(q)\right) \mathrm{d}q \\ + &\int_{(g')^{-1}(1)}^{q^M} \left(\tfrac{1}{2} + g'(q) - c'(q)\right) \mathrm{d}q - \int_{(g')^{-1}(1)}^{q^M} \tfrac{3}{2}\left(\tfrac{1+g'(q)}{2}\right)^2 - c'(q)\, \mathrm{d}q \\ + &\int_{q^M}^{q^\star} \left(\tfrac{1}{2} + g'(q) - c'(q)\right) \mathrm{d}q - 0 \\ = &\; 0 \\ + &\int_{(g')^{-1}(1)}^{q^M} d(q)\, \mathrm{d}q \\ + &\int_{q^M}^{q^\star} \left(\tfrac{1}{2} + g'(q) - c'(q)\right) \mathrm{d}q \end{aligned}$$

which is expression (1.15).

1.5.7 Proof of Proposition 1.4

i) The quality conditional profit reads

$$\Pi^{NS}(q) = \max_\theta \left[g(q) + \theta q\right](1 - \theta)$$

First order condition is

$$\begin{array}{rcl} \theta q - \theta g\left(q\right) - \theta^2 q & \geq & 0 \\ \left[-\left[g\left(q\right) + \theta q\right] + q\right] & \theta \geq & 0 \\ \vartheta\left(q\right) & = & \max\left\{\frac{q-g(q)}{2q}, 0\right\} \end{array}$$

When $\frac{q-g(q)}{2q} = 0$, then everyone is sold good q and $\Pi^{NS}(q) = g(q)$. If instead $\frac{q-g(q)}{2q} > 0$ then only types $\left[\frac{q-g(q)}{2q}, 1\right]$ receive quality q and pay the valuation of the marginal consumer.

We need to show that the screening monopolist stops full bunching before the NS monopolist. The threshold for the screening monopolist $\beta(0)$ solves

$$\frac{1 - g'(q)}{2} = 0 \implies g'(q) = 1$$

The threshold for the NS monopolist $\vartheta^{-1}(0)$ solves

$$\frac{q - g(q)}{2q} = 0 \implies g(q) = q$$

As g is continuous and convex, $\beta(0) < \vartheta^{-1}(0)$ is an immediate implication of the Largange Theorem.

iii) The marginal revenue function $\frac{\mathrm{d}}{\mathrm{d}q}\Pi^{NS}(q)$ is equal to $g'(q)$ when $q - g(q) < 0$ and we have full bunching (serve everyone at price $g(q)$). If $\vartheta(q) \in (0,1)$, we can use the envelope theorem to obtain

$$\begin{array}{rcl} \Pi^{NS}\left(q\right) & = & \max_{\theta}\left[g\left(q\right) + \theta q\right]\left(1 - \theta\right) \\ \frac{\mathrm{d}}{\mathrm{d}q}\Pi^{NS}\left(q\right) & = & \left[g'\left(q\right) + \vartheta\left(q\right)\right]\left(1 - \vartheta\left(q\right)\right) \\ & = & \left[g'\left(q\right) + \frac{q-g(q)}{2q}\right]\left(1 - \frac{q-g(q)}{2q}\right) \\ & = & \frac{(q+g(q))(q-g(q)+2qg'(q))}{4q^2} \end{array}$$

proving the statement.

iii) As in the proof of Proposition 18 *iii*), it will be sufficient to show

$$\frac{\mathrm{d}}{\mathrm{d}q}\Pi^{NS}\left(q\right) \leq V'\left(q\right)$$

strictly when $g'(q) \leq 1$. We distinguish three cases.

If q is such that $g'(q) > 1$, then the constraint is immaterial and marginal revenues coincide.

If q is such that $g'(q) \leq 1$ but $g(q) > q$, then the monopolist screens so its optimal choice is given by

$$V'(q) = \left(\frac{1+g'(q)}{2}\right)^2 > g'(q) = \frac{\mathrm{d}}{\mathrm{d}q}\Pi^{NS}(q)$$

If also $g(q) < q$, then the derivative of the profit function is given by

$$\begin{array}{rcl} \frac{\mathrm{d}}{\mathrm{d}q}\Pi^{NS}(q) & = & \frac{(q+g(q))(q-g(q)+2qg'(q))}{4q^2} \\ & < & \frac{(q+q)(g(q)-g(q)+2qg'(q))}{4q^2} \\ & = & g'(q) \\ & \leq & \left(\frac{1+g'(q)}{2}\right)^2 \\ & = & V'(q) \end{array}$$

where the first inequality uses $g(q) < q$ (twice). This proves in general that $q^{NS} \leq q^M$, strictly whenever $g'(q) \leq 1$ (no full bunching in the monopolist case).

1.5.8 Proof of Proposition 1.5

i) Type θ welfare under the NS monopolist is

$$W^{NS}(\theta) = \begin{cases} \left(\theta - \vartheta\left(q^{NS}\right)\right) q^{NS} & \theta > \vartheta\left(q^{NS}\right) \\ 0 & else \end{cases}$$

a linear function (in θ) which is zero at the exclusion threshold and has slope q^{NS}. If q^M is in region B, then $\vartheta\left(q^{NS}\right) = 0$ and

$$q^M > q^{NS} > \beta(0)$$

As β is continuous in types there exists some $\overline{\theta} > 0$ for which

$$q^{NS} > \beta(\theta) = \rho\left(\theta, q^M\right) \qquad \forall\theta \in \left[0, \overline{\theta}\right]$$

Now pick $\theta \in \left(0, \overline{\theta}\right)$ and notice

$$\begin{array}{rcl} W(\theta) & = & \int_0^\theta \beta(\theta')\,\mathrm{d}\theta' \\ & < & \int_0^\theta q^{NS}\mathrm{d}\theta' \\ & = & \theta q^{NS} \\ & = & W^{NS}(\theta) \end{array}$$

and θ is better-off under the NS policy.

ii) Since in Region B no one is excluded by the NS monopolist, total surplus is

$$\begin{aligned} W^{NS} &= g\left(q^{NS}\right) - c\left(q^{NS}\right) + \tfrac{1}{2}q^{NS} \\ &= \int_0^{q^{NS}} g'\left(q^{NS}\right) - c'\left(q^{NS}\right) + \tfrac{1}{2}\mathrm{d}q \end{aligned}$$

Subtracting this from monopolist surplus 72 and breaking down the integral in regions $\left(g'\right)^{-1}(1) < q^{NS} < q^{M}$ we get

$$\begin{aligned} W^{NS} - W^{M} &= \int_{(g')^{-1}(1)}^{q^{NS}} \left(\tfrac{1}{2} + g'(q) - c'(q)\right)\mathrm{d}q - \int_{(g')^{-1}(1)}^{q^{M}} \tfrac{3}{2}\left(\tfrac{1+g'(q)}{2}\right)^2 - c'(q)\,\mathrm{d}q \\ &+ 0 - \int_{q^{NS}}^{q^{M}} \tfrac{3}{2}\left(\tfrac{1+g'(q)}{2}\right)^2 - c'(q)\,\mathrm{d}q \\ &= \int_{(g')^{-1}(1)}^{q^{NS}} d(q)\,\mathrm{d}q - \int_{q^{NS}}^{q^{M}} \tfrac{3}{2}\left(\tfrac{1+g'(q)}{2}\right)^2 - c'(q)\,\mathrm{d}q \end{aligned}$$

first term are welfare gains from undoing damaging, second are losses from underacquisition (compared to monopolist).

iii) We keep fixed q^{M}. As $c''\left(q^{M}\right) \to \infty$, then $c'\left(q^{M}\right) - c'(q) \to \infty$ for each $q > q^{M}$. As $V'\left(q^{M}\right) - \Pi'\left(q^{M}\right)$ is positive but finite this means $q^{NS} \to q^{M}$ and by ii) above

$$W^{NS} - W^{M} \to \int_{(g')^{-1}(1)}^{q^{M}} d(q)\,\mathrm{d}q$$

when convex costs shut down underacquisition, NS policy in Region B has the only (welfare increasing) effect of undoing screening inefficiencies.

1.5.9 Proof of Proposition 1.6

We avoid order statistics notation and let x, y generic be the realized maximum and second order statistics of the entry vector $\boldsymbol{q}$.

We firstly establish that in any market equilibrium, all qualities $q \leq y$ make zero revenues. Suppose otherwise, that is some firm j makes positive revenues selling a set of qualities bounded above by y. By definition of y at least an active competitor can modify his pricing function to copy the revenue-earner in that quality region, not alter the market pricing function (hence his revenues on other qualities) and share those positive revenues. Without invoking arbitrary tie-breaks, and for future reference, we should notice the competitor can, by quoting a pricing function $p_j(q) - \epsilon$ for ϵ arbitrarily close to 0 appropriate all revenues in the shared region without altering his revenues from other qualities: only marginal types would deviate as $\epsilon \to 0$. As no firm makes revenues on $q \leq y$, it follows $m(q) = 0$ for all $q \leq y$, and in particular $m(y) = 0$.

We now solve the problem of the interim monopolist. Beyond forcing $p_i(y) = 0$ (or, equivalently, the quality allocation to be bounded below by y), competition has no impact on the interim monopolist as feasibility forces competitors $-i$ to quote a price $p_{-i}(q) = \infty$ for all $q > y$.

Using the same steps as for the unconstrained monopolist, we write the problem of the interim monopolist as choosing a type dependent quality allocation rule ρ which is increasing (pairwise comparison of IC) and has image $[y, x]$. The interim monopolist revenues are therefore

$$R_i(x, y) \longrightarrow \max_{\rho(\theta)\in[y,x],\ \text{increasing}} \int_\Theta \left[u(\rho(\theta), \theta) - \rho(\theta)\frac{1 - F(\theta)}{f(\theta)}\right] f(\theta)\,\mathrm{d}\theta$$

Pointwise maximization of the objective delivers the candidate allocation

$$\widetilde{\rho}(x, y, \theta) = \arg\max_{q\in[x,y]} g(q) + q(\theta - h(\theta))$$

From the concavity of the objective first order condition characterizes the interior optimum, hence $\widetilde{\rho}(x, y, \theta) = \beta(\theta)$ if $\beta(\theta) \in [y, x]$. The objective is instead strictly decreasing (on the relevant domain) in q if $\beta(\theta) < y$, strictly increasing if $\beta(\theta) > x$. It therefore follows

$$\widetilde{\rho}(x, y, \theta) = \begin{cases} y & y < \beta(\theta) \\ \beta(\theta) & \beta(\theta) \in [y, x] \\ x & x > \beta(\theta) \end{cases}$$

which is a weakly increasing function (in θ), hence the solution of the interim monopolist problem. This proves the allocation rule (1.21).

We now need to compute the revenue function. Per the discussion above all firms but the interim monopolist make zero revenues. The interim monopolist earns

$$R_i(x, y) = R(y, y) + \int_y^x \frac{\partial}{\partial q} R(q, y)\,\mathrm{d}q \tag{1.51}$$

We have $R(y, y) = 0$ and, given allocation function (1.21), marginal revenues for the interim monopolist coincide with those of the unconstrained monopolist: the marginal quality is assigned to types $[b(q), 1]$ at marginal price $g'(q) + b(q)$. So

$$\frac{\partial}{\partial q} R(q, y) = \begin{cases} 0 & y > q \\ V'(q) & y \leq q \end{cases}$$

And we can rewrite (1.51) as

$$R_i(x,y) \;=\; \int_y^x V'(q)\,\mathrm{d}q \;= V(x) - V(y)$$

Since x, y was generic this proves expression (1.22) (after the appropriate notation changes are made).

1.5.10 Proof of Proposition 1.6

We divide the proof in several steps.

Step 1: **Monopolist is the only equilibrium in pure strategies.**

Consider a generic infinite profile of pure strategies $\widetilde{\boldsymbol{q}}$.[63] Suppose there is i, j with $0 < \widetilde{\boldsymbol{q}}_i \leq \widetilde{\boldsymbol{q}}_j$. Then by (1.22), firm i makes zero revenues but pays positive cost, she is better off staying inactive. So no two firms can choose positive qualities in the pure strategies equilibrium. $\mathbf{0}$ cannot be an equilibrium either as everyone's best response is to produce q^M. $\widetilde{\boldsymbol{q}}$ is therefore a candidate equilibrium in pure strategies only if $\widetilde{\boldsymbol{q}}_i = q^M$ for one and only one i, $\widetilde{\boldsymbol{q}}_{-i} = 0$. We need to show this is indeed an equilibrium. The entrant clearly has no incentive to deviate. Other firms do not enter with a quality below q^M as

$$q \leq q^M \implies R\left(q, q^M\right) = 0 < c(q)$$

If $q > q^M$, the deviator becomes interim monopolist and makes profits

$$\begin{aligned} \Pi(q) &= R\left(q, q^M\right) - c(q) \\ &= V(q) - V\left(q^M\right) - c(q) \\ &= \int_{q^M}^{q} \left(V'(q') - c'(q')\right) \mathrm{d}q' - c\left(q^M\right) \end{aligned}$$

Both summands are negative as as $c'(q) > V'(q)$ above q^M. Now we look for equilibria where there are at least 2 active firms. By Step 1, those equilibria must be in mixed strategies.

Step 2: **Active firms play an atomless distribution with support including 0.**

The support of equilibrium play must contain 0: if the support was bounded below by a strictly positive quality $\underline{q}$, then playing the costly $\underline{q}$ would deliver zero revenues with probability 1, and abstaining is better. Similarly, in equilibrium no firm can choose a quality q with positive probability. If that were the case, all opponents best respond by placing zero probability on an open set including q to get a discrete jump

[63] Where $\widetilde{\boldsymbol{q}}$ is shorthand notation for each firm i plays $\delta_{\widetilde{\boldsymbol{q}}_i}$.

in the probability of winning and (almost) the same profits. But then the firm itself wants to shift the mass away from q, depending on the sign of $V'(q) - c'(q)$.

Step 3: **The distribution of the maximum of opponent's qualities must be** $H(q) = \frac{c'(q)}{V'(q)}$.

We know from (1.22) that the maximum across competitors' realizations x is sufficient to determine firms' revenues. Let $H(x)$ be the distribution of such maximum, which from the previous step we know is continuous on $[0, \overline{q}]$ for some $\overline{q} > 0$. By playing q makes expected profits

$$\Pi(q) = \int_Q R(q,x)\,\mathrm{d}H(x) - c(q)$$

The flat profit condition $\Pi'(q) = 0$, necessary for indifference reads

$$\frac{\partial}{\partial q}\int_Q R(q,x)\,\mathrm{d}H(x) - c'(q) = 0$$

Using Leibnitz rule on an invariant support

$$\frac{\partial}{\partial q}\int_Q R(q,x)\,\mathrm{d}H(x) = \int_Q \frac{\partial}{\partial q}R(q,x)\,\mathrm{d}H(x)$$

Now use again

$$\frac{\partial}{\partial q}R(q,x) = \begin{cases} 0 & x > q \\ V'(q) & x \leq q \end{cases}$$

to write

$$\int_Q \frac{\partial}{\partial q}R(q,x)\,\mathrm{d}H(x) = \int_0^q V'(q)\,\mathrm{d}H(x) = V'(q)\int_0^q \mathrm{d}H(x) = V'(q)\,H(q)$$

Which, substituted in the flat profit condition gives the desideratum

$$H(q) = \frac{c'(q)}{V'(q)}$$

It holds $H(0) = 0$ since, by Proposition 1.1 *ii*), for low q marginal revenues $V'(q)$ is equal to $g'(q)$ approaching ∞ by the Inada condition. H is increasing as c is assumed convex and V is concave. The right extremum of the support is determined by

$$H(q) = 1 \implies c'(q) = V'(q) \implies q = q^M$$

Hence the maximum among $n-1$ competitors is an absolutely continuous random variable with support $\left[0, q^M\right]$.

Step 4: **An equilibrium candidate with n active firms is (1.27)**

The CDF H pins down the distribution of the maximal quality among $n-1$ competitors that makes the n^{th} firm indifferent among any quality $q \in \left[0, q^M\right]$. So for each n we have one (and only one) candidate equilibrium which has everyone plays

$$H_n(q) = [H(q)]^{\frac{1}{n-1}}$$

Also absolutely continuous with support $\left[0, q^M\right]$. This proves expression **(1.27)**. Notice CDF H_n admits a positive density

$$h_n(q) = \frac{1}{n-1}[H(q)]^{\frac{2-n}{n-1}} h(q)$$

It is continuous since $h(q) = \frac{\mathrm{d}}{\mathrm{d}q} H(q)$ is continuous in q in light of continuity of V'' established in Proposition 1.1 *ii*).

Step 5: **Sufficiency**

We are left to prove that this is indeed an equilibrium. All active firms are indifferent across all qualities in $\left[0, q^M\right]$ by the flat profit condition, they are indifferent with abstaining as 0 is in the support of the equilibrium. We are left to prove that firms do not want to produce more than q^M. In that case they would be the interim monopolist for sure, making profits

$$\Pi(q) = \Pi\left(q^M\right) + \int_{q^M}^{q} V'(q') - c'(q')\,\mathrm{d}q'$$

The first summand is zero in expectation by the flat profit condition, while the second term is negative by definition of q^M. That inactive firms do not want to produce any positive quantity is immediate: each of the n firms, competing against $n-1$ opponents makes zero profits in expectation and competing competing against n firms increases (in the sense of FOSD) the distribution of the best competitors' quality. Inactive firms are strictly better off abstaining completing the proof that this is an equilibrium.

1.5.11 Proof of Theorem 1.2

i) We preliminary derive type dependent consumer for each realized competitive environment x, y adding to the utility of the lowest type the integral of allocations (equal to u_θ) characterized in Proposition 1.6

$$\begin{aligned} W(\theta,x,y) &= u(0,x,y)+\int_0^\theta u_\theta(\rho(\theta',x,y),\theta')\,\mathrm{d}\theta' \\ &= g(y)+\int_0^\theta \max\{y,\min\{x,\beta(\theta')\}\}\,\mathrm{d}\theta' \end{aligned}$$

which is expression (1.29) in the main text (and for $x=q$, $y=0$ give surplus under monopoly). It is immediate to notice that type dependent welfare is increasing in (x,y)

$$\forall\theta,\quad (x,y)\geq_2(x',y') \text{ implies } W(\theta,(x,y))\geq W(\theta,(x,y'))$$

where $\geq_2$ is the standard incomplete order in $\mathbb{R}^2$. Given monotonicity of value conditional on the realized qualities, to establish the result is sufficient to show that the random vector of marketed qualities $\boldsymbol{x},\boldsymbol{y}$ has distribution ordered according to first order stochastic dominance (FOSD) in the equilibria with active competition.

Using individual firms' equilibrium play (1.27) we derive the distribution of the maximal quality $\boldsymbol{x}$ in equilibria with n active firms

$$\begin{aligned} H_n[x] &= Pr[\max\{q_1,q_2,\ldots q_n\}\leq x] \\ &= \prod_{i=1}^n[q_i\leq x] \\ &= [H_n(x)]^n \\ &= [H(x)]^{\frac{n}{n-1}} \end{aligned}$$

Since

$$s(n)=v^{\frac{n}{n-1}}$$

is a (strictly) increasing function of n for every $v\in[0,1]$, it follows

$$n>m \implies H_n[x]>H_m[x]\ \forall x$$

the best quality in equilibria with lower competition first order stochastically dominates the best quality in equilibria with more intense competition. Now we use the following fact[64]

Fact. *Let $X_1,\ldots X_n$ be independent observations from a continuous CDF F. Then, the conditional distribution of the second order statistic given $\max_{i\in[n]}X_i=x$ is the same as the unconditional distribution of the maximum in a sample of size $n-1$ from a new distribution, namely the original F truncated at the right at x.*

[64]See theorem 6.7 in http://www.stat.purdue.edu/~dasgupta/orderstats.pdf.

It follows from the fact that the distribution of each other firm's quality conditional on $\boldsymbol{x} = x$ in an equilibrium with n active firms is

$$H_{x,n}(q) = \left[\frac{H(q)}{H(x)}\right]^{\frac{1}{n-1}} \mathbb{I}\{q \in [0, x]\}$$

so $\boldsymbol{y}\,|x$, the maximum across $n-1$ of them is distributed

$$H_{x,n}(y) = \left[\left[\frac{H(y)}{H(x)}\right]^{\frac{1}{n-1}}\right]^{n-1} \mathbb{I}\{y \in [0, x]\} = \frac{H(y)}{H(x)}\mathbb{I}\{y \in [0, x]\}$$

In particular, is independent of n. As the distribution of $\boldsymbol{x}$ is FOSD ranked across equilibria and the distribution of $\boldsymbol{y}$ given $\boldsymbol{x}$ is invariant across equilibria, it follows the joint of $\boldsymbol{x}, \boldsymbol{y}$ is FOSD across equilibria, which we argued above is sufficient for the statement.

ii) Suppose $y \leq x \leq q^M \leq (g')^{-1}(1)$. As both market statistics realize in the full bunching region, the competitive allocation (1.21) will assign every type the undamaged quality x at price $g(x) - g(y)$. Per (1.29), type dependent surplus is

$$W(x, y, \theta) = g(y) + \theta x$$

and total surplus is

$$W(x, y) = \int_{\Theta} W(x, y, \theta)\, \mathrm{d}\theta = g(y) + \frac{1}{2}x$$

So

$$W_n = \mathbb{E}_n[W(\boldsymbol{x}, \boldsymbol{y})] = \mathbb{E}_x\left[\mathbb{E}_{y|x}[g(\boldsymbol{y})] + \frac{1}{2}\boldsymbol{x}\right] \tag{1.52}$$

Since

$$H_x(y) = \frac{g'(x)}{c'(x)}\frac{c'(y)}{g'(y)}, \quad y \in [0, x]$$

is the conditional CDF of $\boldsymbol{y}$ given $\boldsymbol{x} = x$ it follows

$$\begin{aligned}
\mathbb{E}_{y|x}[g(\boldsymbol{y})] &= \int_0^x g(y)\, \mathrm{d}H_x(y) \\
&= \tfrac{1}{H(x)}\left(H(y)\,g(y)\Big|_0^x - \int_0^x g'(y)\,H(y)\,\mathrm{d}y\right) \\
&= \tfrac{1}{H(x)}\left[(H(x)\,g(x)) - \int_0^x g'(y)\,\tfrac{c'(y)}{g'(y)}\mathrm{d}y\right] \\
&= g(x) - \tfrac{c(x)}{H(x)} \\
&= g(x) - c(x)\,\tfrac{g'(x)}{c'(x)}
\end{aligned} \tag{1.53}$$

The function

$$s(x) = g(x) - c(x)\frac{g'(x)}{c'(x)}$$

has $s(0) = 0$ and

$$\begin{aligned} s'(x) &= g'(x) - c'(x)\tfrac{g'(x)}{c'(x)} - c(x)\tfrac{g''(x)c'(x) - g'(x)c''(x)}{[c'](x)^2} \\ &= -c(x)\tfrac{g''(x)c'(x) - g'(x)c''(x)}{[c'](x)^2} \\ &> 0 \end{aligned}$$

so s is positive and monotonically increasing in $\left[0, q^M\right]$, and $s\left(q^M\right) = g\left(q^M\right) - c\left(q^M\right)$

Now, we substitute (1.53) into (1.52) to get

$$\begin{aligned} W_n &= \mathbb{E}_n\left[g(\boldsymbol{x}) - c(\boldsymbol{x})\tfrac{g'(\boldsymbol{x})}{c'(\boldsymbol{x})} + \tfrac{1}{2}\boldsymbol{x}\right] \\ &= \int_0^{q^M}\left[g(x) - c(x)\tfrac{g'(x)}{c'(x)} + \tfrac{1}{2}x\right]\mathrm{d}H^{\frac{n}{n-1}}(x) \end{aligned}$$

Integrating by parts the last expression we obtain

$$\left[g\left(q^M\right) - c\left(q^M\right) + \frac{1}{2}q^M\right] - \int_0^{q^M}\left[g'(x) - \frac{\mathrm{d}}{\mathrm{d}x}\left(c(x)\frac{g'(x)}{c'(x)}\right) + \frac{1}{2}\right]H^{\frac{n}{n-1}}(x)\,\mathrm{d}x$$

the first term is monopolist profit, while we can notice

$$g'(x) - \frac{\mathrm{d}}{\mathrm{d}x}\left(c(x)\frac{g'(x)}{c'(x)}\right) = s'(x) > 0$$

as proved above, so we are subtracting the integral of a positive function, reducing monopolist welfare and proving the statement.

We now prove the general welfare decomposition (1.32). Fix a realization of market statistics x, y. We can write

$$W(x,y) = \int_0^x \frac{\mathrm{d}}{\mathrm{d}q}W(x,y)(q)\,\mathrm{d}q$$

where $\frac{\mathrm{d}}{\mathrm{d}q}W(x,y)(q)$ is the marginal contribution to welfare . It holds

$$\frac{\mathrm{d}}{\mathrm{d}q}W(x,y)(q) = \begin{cases} \left(S^{FB}\right)'(q) - c'(q) & q < y \\ m(q) & q \in [y, x] \end{cases}$$

since for qualities below y welfare from competition grows as first best (since all all goods are distributed to everyone) net of the additional marginal costs (incurred twice). In the quality region $[y, x]$ surplus grows as monopolist since those qualities are assigned to the same types and cost incurred only once. It then follows

$$W(x,y) - W_M = \int_0^x \frac{\mathrm{d}}{\mathrm{d}q} W(x,y)(q)\,\mathrm{d}q - \int_0^{q^M} (W_M)'(q)\,\mathrm{d}q =$$

As monopolist surplus grows as first best below $(g')^{-1}$ and as the second branch of $m(q)$ in $\left[(g')^{-1}, q^M\right]$ we rewrite the difference (assuming $y > (g')^{-1}$)

$$= -c(y) + \underbrace{\int_0^{(g')^{-1}} \left[\left(S^{FB}\right)'(q) - \left(S^{FB}\right)'(q)\right]\mathrm{d}q}_{=0} + \int_{(g')^{-1}}^{y} \underbrace{\left[\left(S^{FB}\right)'(q) - (W_M)'(q)\right]}_{=d(q)}\mathrm{d}q$$

$$+ \underbrace{\int_y^x m(q) - m(q)\,\mathrm{d}q}_{=0} - \int_x^{q^M} m(q)\,\mathrm{d}q$$

$$= \int_{(g')^{-1}}^{y} d(q)\,\mathrm{d}q - c(y) - \int_x^{q^M} m(q)\,\mathrm{d}q$$

Taking expectations under the duopoly equilibrium distribution of market statistics $\boldsymbol{x}, \boldsymbol{y}$,

$$W_2 = \mathbb{E}_2\left[W(\boldsymbol{x}, \boldsymbol{y})\right]$$

we obtain the fundamental welfare decomposition

$$W_2 - W_M = \mathbb{E}_2\left[\int_{(g')^{-1}(1)}^{\boldsymbol{y}} d(q)\,\mathrm{d}q - c(\boldsymbol{y})\right] - \mathbb{E}_2\left[\int_{\boldsymbol{x}}^{q^M} \left[m(q) - c'(q)\right]\mathrm{d}q\right]$$

where we mean

$$\int_{(g')^{-1}(1)}^{y} d(q)\,\mathrm{d}q = 0$$

whenever $y < (g')^{-1}(1)$.

We are now left to show that under convex cost $c(q) = q^\alpha$ as α grows to infinity equilibria with active competition dominate monopoly. As $q_M(\alpha)$ solves $V'(q) = \alpha q^{\alpha-1}$, irrespectively of the revenue function monopolist quality will converge to the point at which the (marginal) cost function explodes, that is

$$q_{\infty}^{M}=\lim_{\alpha\to\infty}q^{M}\left(\alpha\right)=1$$

while its cost will converge to zero

$$c_{\infty}^{M}=\lim_{\alpha\to\infty}\left(q^{M}\left(\alpha\right)\right)^{\alpha}=0$$

also, since

$$\lim_{\alpha\to\infty}c_{\alpha}'\left(x\right)\to\begin{cases}0 & x<1\\ 1 & x=1\\ \infty & x>1\end{cases}$$

it follows that

$$H_{2,\alpha}\left(q\right)=\begin{cases}\frac{c'(q)}{V'(q)} & q\leq q^{M}\left(\alpha\right)\\ 1 & q>q^{M}\left(\alpha\right)\end{cases}\longrightarrow\begin{cases}0 & q<q_{\infty}^{M}\\ 1 & q\geq q_{\infty}^{M}\end{cases}$$

which means that each firm's equilibrium strategy converges in probability to q_{∞}^{M}. Also, notice that as $\boldsymbol{y}\leq q^{M}\left(\alpha\right)$ then

$$\mathbb{E}_{\alpha}\left[c\left(\boldsymbol{y}\right)\right]\leq c\left(q_{\alpha}^{M}\right)\to0$$

Plugging those results in (1.32) we observe that the limit welfare impact of competition is given by

$$\begin{aligned}&\int_{(g')^{-1}(1)}^{q_{\infty}^{M}}d\left(q\right)\mathrm{d}q-c_{\infty}^{M}-\int_{q_{\infty}^{M}}^{q_{\infty}^{M}}\left[m\left(q\right)-c'\left(q\right)\right]\mathrm{d}q\\ &=\int_{(g')^{-1}(1)}^{q_{\infty}^{M}}d\left(q\right)\mathrm{d}q \qquad (1.54)\end{aligned}$$

proving the statement.

1.5.12 Proof of Proposition 1.8

i) By the same reasoning as in the proof of Proposition 1.7, one firm offering the monopolist pricing function $p^{M}\left(\cdot\right)$ is the only candidate equilibrium: no two firms can quote a non-trivial pricing function but someone must. However in this case, an inactive firm can offer pricing function $p^M(\cdot) - \epsilon$ (so that allocations would be unchanged), pay $c\left(q^M\right)$ and make revenues that are $\epsilon-$close to $V\left(q^M\right)$, a profitable deviation. There is no RS equilibrium.

For the second statement in *ii*) we notice that full interim stability of the monopolist follows from Proposition 1.7 *i*) (it is a Nash equilibrium of the first stage game), while its ex-post instability is proved above.

To prove the remaining part of *ii*) and *iii*) we preliminarily notice that

$$0 = R\left(q^M, q^M\right) < c\left(q^M\right) < R\left(q^M, 0\right) = V\left(q^M\right)$$

and $R\left(q^M, x\right)$ is monotonically decreasing, so there will be a threshold $m^\star$ at which $R\left(m^\star, q^M\right) = c\left(q^M\right)$. By Proposition 1.7 a potential entrant against best quality x can make revenues

$$\widetilde{R}(x) = \max_q R(q, x) - c(q)$$

As $R(x, q) = V(q) - V(x)$, x induces a sunk cost, upon entry q^M is played and

$$\widetilde{R}(x) = \max\left\{V\left(q^M\right) - V(x) - c\left(q^M\right)\right\}$$

From which it follows

$$\widetilde{R}(x) > 0 \iff x < m^\star$$

By a similar argument, an ex-post deviation occurs if and only if the second quality is above $m^\star$. A deviator that enters with q^M can indeed offer allocation $\rho\left(y, q^M, \cdot\right)$ and make revenues that are arbitrarily close to $V\left(q^M\right) - V(y)$: to do that he must grant a small discount to all types in $[b(y), b(x)]$ so to push the interim monopolist out of the market. Whenever y realizes strictly below $m^\star$ therefore an ex-post entrant the profitable deviation of offering the interim-contract $\left[y, q^M\right]$, with prices ϵ-reduced to all types for which $\beta(\theta)$ is below the level owned by the realized interim monopolist.

It therefore follows that the degree of interim stability of the $n-$equilibrium is $H_{x,n}(m^\star)$ and the degree of ex-post stability is $H_{y,n}(m^\star)$, where $H_{x,n}$, $H_{y,n}$ are the CDF of market statistics calculated under equilibrium play with n active firms. Point *ii*) now follows from $\boldsymbol{x} \geq \boldsymbol{y}$ by definition, while the $n-$ranking in *iii*) follows from the FOSD ranking of market statistics proved in Proposition 1.2 *i*).

Chapter 2

Price Incentivation Into Social Programs: Estimation and Control

2.1 Introduction

2.1.1 Motivation

It has become commonly accepted that policy interventions should be data-driven, and that decisions about enactment of a certain social program should be based on empirical evidence. However, it is complicated to argue what dataset and what type of statistical analysis is relevant to provide the empirical evidence to inform a certain decision, especially if the latter is a relatively new intervention. Clearly, the best way to know the effectiveness of a certain policy is to enact the policy itself. However, once the policy is enacted, in a one-shot problem we are left with no informed decisions to make: the dataset is in some sense useless. If the problem is not one shot, and on the contrary we expect to face the same problem in subsequent periods, then the dataset so collected provides valuable information about the environment the decision maker is going to operate in the future. Moreover, if the planner can choose among a set of "similar" policies, say she can vary the level on coverage and the intensity of the monetary incentive, then the return of a given policy can, through an economic theory, provide information about the likely returns of the unchosen alternatives. Motivated by this observation, we now present a stylized example of problems that will be studied in this chapter. It is kept simple enough to be numerically solvable (a task we accomplish in Section 3) at the cost of shutting down most of the interesting channels that will be explored in details in Sections 4 and subsequent.

Example 2.1. Consider the problem of the social planner of an island where in each period a finite number N of citizens spend their whole life. At time 0 new disease appears on the island, which would make at each future period each current inhabitant ill with probability $\omega \in \Omega = [0,1]$, unknown but constant through time. There is a vaccine for the disease, that the planner can purchase at unit cost q^* and that would make the recipient immune to the disease. The planner must in each period commit to a price q at which citizens can purchase the vaccine; she has to provide vaccine to whomever agrees to pay the fixed price, which results in total social (net) cost $C(q) = D(q)(q^* - q)$. The demand schedule is known, and it is known that determinants of demand are uncorrelated with the likelihood of being ill after no treatment. The planner observes $s_t \in (\{0,1\} \times \{0,1\})^N$ containing the treatment choice and health outcome pairs for every citizen in the population, and uses health statuses for individuals that chose not to vaccinate (that get ill with probability ω) to update his belief $p \in \Delta(\Omega)$. Social utility depends on both the distribution of health statuses across the population and the monetary (net) profit from selling the vaccine.

2.1.2 Policymaking as a MultiArmed Bandit problem

A MultiArmed Bandit (MAB) problem is an infinitely repeated decision problem where the decision maker chooses at each period to pull one among several arms that yield stochastic returns. The outcome from each period enters current utility and is used to estimate the return of the arm that is pulled and, in case returns of arms are believed to be correlated, of the whole set of arms.

In his exposition of the econometric approach to causal inference, Heckman (2008) proposes a three step procedure which resemble a model of policymaking as a MAB where arms are treatment incentive schemes whose payoff value and correlation is disciplined by an economic theory represented by a reduced form model that is estimated as the decision process unfolds.

The aim of this chapter is to provide a formal model of policymaking as the decision problem of an incentive designer that solves the following MAB: the state space contains models of individual decision making and a distribution over individual types that determine stimulus response and post-treatment outcomes,[1] while the action set is composed of a stimulus dimension and a dimension which determines the data available at the end of each period for estimation. Stationarity is imposed in

[1] So it is not defined before the set of stimuli (policies) that can be implemented is given.

the form of having neither the unknown social state nor the set of incentives available to vary across periods. Data collected in each period provide a signal of the social state that is imperfect for two reasons: firstly the current population (collection of citizens alive) is only a finite sequence of independent random draws from the invariant super-population,[2] secondly the social state contains information about all possible policies, and it may be possible that the one currently chosen is not able to identify some dimension of the social state.[3] Notice that, although the general structure is essentially that of a bandit problem, both the state space and the action space are unusually rich for a MAB setup. One of the contributions of this chapter is to offer this representation. We offer sufficient condition on individual preferences that are sufficient to implement a rich class of monetary incentivation schemes, the Becker-DeGroot-Marschack (Becker et al. (1964)) meachinism with treatment lotteries as prices. This set of mechanisms includes as extreme points fully coercive (RCT) and fully voluntary (posted price) mechanism, but allow for intermediate incentivation through stochastic assignment.

Framing this type of problem as an infinitely repeated statistical game of the type studied in part 2 has advantages and disadvantages. A clear advantage is that, by giving the planner a decision problem it implicitly defines the parameters that are relevant for policymaking. This is not a trivial point, and has risen prominent critiques on the literature of treatment effect estimation. In Heckman (2008) words, "what is often missing in the literature of treatment effect estimation is a clear discussion of the policy question being addressed by the particular treatment effect being identified and why it is interesting". A similar message is given by Deaton (2010) "I shall argue that the analysis of projects needs to be refocused towards the investigation of potentially generalizable mechanisms that explain why and in what contexts projects can be expected to work". Moreover, the provides the planner with some instruments to generate datasets that identify and estimate those parameters. Also, contrary to static information acquisition problems, it makes *indirect costs* of experimentation appear, associated to the static suboptimality of most informative mechanisms.[4] An equally clear disadvantage is that stationarity gives too much rigidity and precludes modifications in the DGP and in incentivation tools that may

[2]The name, as well as most terminology, is taken from Imbens and Rubin.

[3]Once we look at policymaking as exploration in a bandit problem with correlated arms, policies have an experimental dimension as they produce a dataset that can be used to estimate the social state and therefore the return of alternative policies.

[4]Without a framework that disciplines dynamic decision making and estimation possibilities the study of this trade-off would be impossible, although it seems that indirect cost from distorting the assignment mechanism may be relevant in particular applications.

instead occur in reality.[5] Having the problem repeated may also suggest that the estimation problem is not really about causal inference but is simple prediction: even without a model of individual decision making the planner can estimate the distribution of social outcomes conditional on a stimulus to estimate the "return" of that stimulus. Without an economic theory the planner would be playing a bandit problem with uncorrelated arms: after choosing a stimulus he still gets to estimate its returns, but this does not inform him about any counterfactual policy. Prior conjecturing exercise gives likelihoods to different models of individual behavior that disciplines the correlation across returns on stimuli. Inference on the parameters of those model is used to inform about returns of (potentially) all the arms.

The class of problems that we aim to study is categorized by Heckman and Vytlacil (2007) as $P3$: *Forecasting the impacts of interventions never historically experienced to various environments, including their impacts in terms of well-being.* A new problem appears (the illness appears on the island in the example above) or new incentive schemes becomes possible (subsidy to attend school, forcing unemployed to participate into labor training programs...); planner has to act in a "novel" environment the he think is going to remain reasonably stable for a reasonably long period of time. In each of those periods he can provide citizens a (monetary) incentive that affect individual selection into programs as well as the (distribution over) datasets that are used in the sequential estimation of the social state.

2.1.3 Literature Review

We now proceed reviewing the literature on multiarmed bandit problems and their applications in economics, as well as earlier contributions that at firstly using selection choices as a source of information and secondly designing selection mechanism so they provide better information.

Bandit Problems in Economics

Multiarmed bandit problems have found various applications in the economics literature, starting from the pioneering work of Wald (1947, 1973) on the sequential design on optimal statistical test. Wald's theory was then extended applied to the framweork of a "statistical game" by Blackwell and Girshick (1979); Blackwell (1951). Easley and Kiefer (1988) present the recursive characterization of the multiarmed bandit problem that allows for infinite and correlated arms;[6] our general exposition in

[5]As I discussed in the introduction, I want to interpret infinite repetition in a stationary environment as a modeling device by the planner who is confident the social environment will remain stable over a long enough period.

[6]They claim their setup differs from that of the bandit problem as they have continuous action space. It seems that, although important, such generalization would not per se represent a crucial departure from the standard bandit problem. I find the real novelty of their approach to be the

Section 2 heavily builds on their contribution. Aghion et al. (1991) give an alternative derivation of their result under slightly more general assumptions and then analyze the asymptotic properties of those learning problems, with particular attention on conditions under which "adequate learning" occurs that allows the decision maker to choose asymptotically the correct action (best response to the true state). Both papers focus on charactering the joint properties of actions that are played infinitely often and the limit belief. They prove[7] that actions played infinitely often must be a static best response to the limit belief, and all states in the support of the limit distribution induce (jointly with each limit action) the same distribution over signals as the one generated by the true state. Those joint properties of limit actions and limit belief characterize the concept of self-confirming equilibrium (SCE) in (static) strategic interaction settings.[8] Study of the asymptotic behavior of such problem is important as it permits to address conditions under which adequate learning occurs and to give a foundation to the SCE concept, but there is no reason to disregard the transition dynamics as uninteresting. The approach we take in this chapter is somehow opposite to that of the SCE literature: rather than founding a static equilibrium concept as the rest point of an unmodeled learning process, we study the decision problem at early stages of experimentation and use the infinite repetition of the game as an approximate modeling device that the agent adopts to describe what are his possibilities to learn and employ the acquired information.[9] Bandit problems have been applied to economic problems since Rothschild (1974), who studied the behavior of rational and optimizing sellers in a market where they are initially ignorant of the demand curves $D(p, \theta)$, $\theta \in \Theta$ they face; it is the first well-known explicit formulation of an economic problem in terms the bandit problem. Bergemann and Välimäki (1996) study the other side of the market, where consumers address from different sellers that produce goods for which they have tastes that are ex-ante unknown but can be experimented. Moscarini (2005) studies a labor market in which employer learns progressively the quality of the match and may decide to terminate the relationship. Bolton and Harris (1999) study strategic interaction into exploration, addressing the issue of information externalities when

conceptual distinction between the estimation and control properties of an action, which is even cleaner in Aghion et al. (1991) who abandon the necessity of observable utility.

[7]Lemma 3 and 4 in Easley and Kiefer (1988), Theorem 2.2 and 2.4 in Aghion et al. (1991).

[8]For an analysis of how the SCE concept has originated and his relevance for the analysis of adaptive processes in a repeated iteration context see the survey by Battigalli et al. (1992).

[9]Such approximation is valid in case the agent is confident that the problem will not change substantially in the near future and has a high enough discount factor to make asymptotic outcomes (hence in regions where it is more difficult to justify the stationarity assumption) utility-negligible.

experimentation results from one agent can be used to by other agents. We pointed at different branches of the economics literature where the MAB was applied to show how pervasive the exploration-exploitation trade-off is; for a reader that is interested in a more in-depth review, Bergemann and Valimaki (2006) offer an excellent survey of bandit problems in economics and finance.

The Econometric approach to Causal Inference

It is somehow surprising, since the prototypical bandit model of Thompson (1933) was introduced to study clinical trials, that advancements in research on social program evaluations (from the pathbreaking contribution of Heckman (1977)) have not until recently stimulated extensions of the exploration-exploitation problem to this field. That is the aim of this Chapter.

A social program offers the population (citizens alive at a certain period) a set of mutually exclusive statuses: being vaccinated against a disease, obtaining (higher) education, participating into labor training programs are only some examples. Participation is not necessarily mandatory, but is the result of an individual selection choice. Modeling the individual selection problem as a source of information about structural parameters rather than a source of concern for failures of statistical assumptions upon which the properties of estimators are based (as is the case for the statistical approach, see Holland (1986)), is the distinguishing feature of the econometric approach to treatment effect estimation, pioneered by Heckman (1977). Its exposition is detailed in the Handbook of Econometrics chapter by Heckman and Vytlacil (2007). To use a language that conforms with the literature on estimation of treatment effects we will be as adherent as possible to the terminology used in Imbens and Rubin (2015) in terms of statistical assumptions (SUTVA, unconfoundedness, etc.), and objects of interest (superpopulation, propensity score).

Before choosing an incentive the planner must form a conjecture about a fairly complicated object as a model for selection and post-treatment choice. In his exposition of the econometric approach to causal inference, Heckman (2008) proposes a **three step procedure** that to "construct" such conjecture. The first stage has the researcher specify a scientific/economic theory describing a mapping from unknown social parameters to stimuli-contingent social outcome distributions, that is he must define the (incentive design) decision problem. It is clear that different incentive schemes require knowledge of different individual characteristics.[10] The second and third stages in the procedure make use of a dataset that is assumed exogenously

[10]Suppose planner can run an information campaign about efficacy of a program: relevant parameters would be individual prior beliefs and how he reacts to information (attention, trust on information etc). Conversely if he considers selling the treatment, then relevant parameters will

available to make inference about parameters that are defined relevant by the first stage. The second stage addresses concerns of identification, that is whether the researcher can recover the parameters from the population from which the dataset is drawn. The third step evaluates how the planner can use the available finite sample to obtain the most efficient estimators of the parameters, giving confidence interval on such estimates that allow to test hypothesis and make informed decisions. Identification and estimation issues are addressed by conjecturing the process that generated the data available (DGP). Beyond the distribution of individual characteristics, this DGP *includes the institutional settings that incentivized choice in previous periods.* The econometrician following the procedure described above defines relevant parameters by formulating a future policy problem, and infers parameters that are relevant for this problem by using an exogenously available dataset to estimate a DGP that contains a description of the policies enacted in the past.[11]

The aim of this chapter is to study the problem of an incentive designer that solves a multiarmed bandit problem where

- The **state space** contains models of individual decision making and a distribution over individual types that determine stimulus response and post-treatment outcomes,[12] and
- The **action set** is composed of a stimulus dimension and a dimension which determines the data available at the end of each period for estimation.

Stationarity is imposed in the form of having neither the unknown social state nor the set of incentives available to vary across periods. Notice that under this formulation we change the nature of phases 2 and 3 of the econometric procedure, from a purely conjectural exercise to an information acquisition choice: when acting inside a period decision problem (stage 1) the policymaker takes into account of his future selves that will solve the same decision problem using information contained in a dataset collected (also) in the current period and that parameters identified in this sample (and the quality of their estimators) depend on the incentive mechanism itself.

be the all the determinants of the value of treatment that is compared to the price when making the decision, joint with income constraints, and others.

[11] Even if those policies may be different from those that he is considering to implement.

[12] So it is not defined before the set of stimuli (policies) that can be implemented is given.

Experimental design and the CPS contribution

The economteric apprach to causal inference recognizes the selection assignment mechanism as a key source of information. Chassang et al. (2012) (CPS from now on) discuss how an experimenter can use treatment assignment mechanisms to elicit individual characteristics that are relevant for policymaking. Their innovative idea is to use messages sent in a direct treatment assignment mechanisms in order to estimate the distribution of selection-relevant types.[13] One of the selection mechanism they propose[14] makes use of the **Becker-DeGroot-Marschack** (Becker et al. (1964)) mechanism to elicit under weak assumptions the individual reservation price for treatment. After running a "rich enough" BDM mechanism on a population of infinitely many (continuum) individuals, planner learns the distribution of the selection variable and, provided a positive fraction of each type is treated, the type-conditional potential outcome. What is crucial is that this estimation informs about the return of *all other BDM policies*, which of course include much different incentive schemes than that runned at the experimenting stage. We take their intuition that the selection mechanism is (also) a source of information and use BDM assignment mechanisms in an environment that is different since

1. Experimentation occurs in "real time" rather than at an ex-ante information acquisition stage. When experimentation (in particular this type of social experimentation) and decision times differ, one should be clear on how the ex-ante stage is different from the decision stage: why *should not* we care about realizations at this stage in our welfare calculations[15] and why, on the contrary, *should* we care about the ex-ante stage, that is how data collected in this stage inform us about future policymaking (inter-temporal external validity)? This requires both a description of the decision problem to be faced and assumptions on how the population in the experimentation phase is related to that of the decision problem.

[13] As the treatment status is binary in order to create an environment that can discriminate along a reasonably rich set of characteristic they need to consider stochastic treatment assignment mechanisms.

[14] They also consider more complicated mechanism that elicit other individual characteristics (belief over the return of the treatment technology, etc), but in this paper I will focus only on reservation price conditional (BDM) incentivation.

[15] There are two possibilities for this to happen. Either the planner does not discount the future so that only asymptotic results are welfare relevant and all periods are information acquisition periods, or the effective population over which experimentation is carried is so small not to impact the aggregate social welfare. Under the latter case concerns of external validity would emerge (see subsection), beyond making the infinite experimentation unit assumption even more critical.

2. Without specifying a "technology" to use information we cannot compare the information accuracy of a selection mechanism with the direct cost (pay the researcher, participation fees, data elaboration...) of running such mechanism. If experimentation occurred in real time, beyond direct costs we would also have to care about indirect cost, that is suboptimality in terms of expected current reward of the most informative mechanisms.[16] Absence a specified decision problem does not undermine validity of the CPS approach, as they rank mechanisms according to the Blackwell order.[17] However, the Blackwell order is an incomplete order and we miss quantitative statements about the value of estimating a certain parameter.

3. Focusing on identification issues by assuming a continuum of individuals misses important issues related to estimation, especially in case we are estimating nonparamertically a density of reservation prices and, conditional on each such price, a pair of potential outcomes. In the vaccination example we solve in Section 3 as a motivating example, we show that focusing on identification problem only (that is, letting $N \to \infty$) makes the exploration-exploitation trade-off disappear.

2.1.4 Plan of the Chapter and Contributions

The Chapter proceeds as follows. Section 2 gives the formal setup of the multi-armed bandit problem and characterizes his solution in the recursive form as in Easley and Kiefer (1988) and Aghion et al. (1991). As an original contribution, we show how to associate each MAB to an uncertainty function (in the sense of DeGroot (1962)) so that the implied information function is traded-off one for one with expected utility at each belief state to determine the optimal policy. In principle this uncertainty measure would identify the set of relevant parameters and quantify the estimation content of selection mechanisms. However, since the state space in the model of policymaking we build is usually very rich and the measure is an endogenous object, this characterization may be of limited partical use. In Section 3 we numerically solve the vaccination example presented in the introduction to show one instance the

[16]To clarify this point, suppose that high reservation price was associated to low efficacy of treatment: as incentive compatibility requires likelihood of being treated to increase in reservation price (Proposition 2.3), a perfectly informative mechanism would induce adverse selection.

[17]Hence, by the Blackwell (1951) theorem we know that information superior mechanism will be welfare improving irrespectively of the particular problem the planner will face.

uncertainty function is tractable and how it informs about the estimation-exploitation trade-off in a simple incentive selection problem; the example also provides a useful benchmark with which one can compare the more general assumption maintained later. Most novel contribution are contained in Section 4 and subsequent, where we study price incentivation into social programs as MAB problem.

Section 2.4.2 details the derivation of the class of reduced form model over which beliefs evolve along the experimentation path. The action space is instead constituted of all possible stimuli government can provide to alter the individual pre-selection problem and by monitoring decisions that will determine the dataset available at the end of the period. Section 2.4.3 presents a class of stimuli (BDM mechanisms) that induce different response along a (single dimensional) reservation price type. A selection intensity maps reduced form models into the proportion of individuals that choose treatment (propensity score) and the outcome of treatment (success rate). A regime is defined as a class of selection intensities that are ranked according to an intensity index (higher intensity selections have each type participate to treatment with higher probability), and rich enough to have for each nondegenerate model a map from propensity score to elements of the class which induce them. Selections inside the same regime share similar incentive structure: RCT (fully coercive) and posted price mechanisms (fully voluntary) are examples. Section 2.5 discusses how to validly implement a stimulus as a reservation price conditional selection intensity inside the finite population. A sampling procedure is a map that takes a pair of population and assignment intensity and returns a joint probability distribution over sampling outcomes. It is ϵ valid if for each pair selection mechanism and social model the expected sample size is ϵ close to the superpopulation propensity score. Bernuolli trial is the simplest example of a valid sampling procedure, but it induces the largest variance of the realized propensity score. We propose a sampling procedure that exploits correlation in individual assignment to minimize the variance of the propensity score (which is desirable, Kasy (2013)) subjet to the constraint that the sampling intensity is validly implemented. Also, implicit implementation of the BDM mechanism preserves information continuity of the selection intensity, contrary to models in which only the identification domain is of interest (Chassang et al. (2012)). Section 2.6 finally studies the outcome model, namely the post-selection determinant of individual "health" outcomes. We show how each regime is characterized by a novel object, the *distortion function* which measures the excess success rate of a particular selection intensity compared to the coercive regime (RCT) that achieves the same propensity score. We show that the unconfoundedness assumption, which in this specification is a *joint* statement about the reduced form model and the selection mechanism, is *characterized* by a unitary distortion function. We also prove that, under linear social preferences, fully voluntary (posted price) mechanisms are control optimal, that is the static policy function always intersect the conditionally deterministic regime. As valid implementation of posted price mechanisms induces the largest variance of the realized propensity score among regimes in the BDM class, their control optimality presents a novel estimation-eploitation trade-off. To summarize, Sections 4, 5 and 6 have three main contributions:

1. We provide a modeling contribution: by offering a formal model of policymaking as a MAB we give a unifying interpretation of the econometric approach to causal inference and experimental design. As a byproduct of this modeling exercise, we introduce the novel concepts of *regimes* as a collection of stimuli that share the same incentivation scheme (ranging from fully coercive to fully voluntary incentivations), and of *regime distortion functions* that naturally map into perturbations of the unconfoundedness assumption.

2. We provide (weak) sufficient conditions for implementation of BDM mechanisms. Using evidence obtained in simpler settings, we end up quesitoning the validity of the BDM in this context in light of the completixy of the treatment lottery. Disregarding those concerns, we propose a sampling procedure that ($\epsilon-$)validly implements all BDM mechanisms, while *i*) preserving information continuity (contrary to the CPS framework) and *ii*) minimizing the variance of the propensity score among valid procedures.

3. We highlight a novel channel through which the control optimality conflicts with optimal information acquisition: posted price price mechanisms are control optimal (under linear social preferences), but also induce the largest variance (among BDM mechanisms) of the empirical propensity score under the proposed sampling procedure.

2.2 Multiarmed Bandit Problems and DeGroot Uncertainty Functions

In this section we setup the multiarmed bandit problem and introduce some notation that will be used throughout the Chapter. We characterize each action by its material consequences (Anscombe and Aumann act) and by its estimation properties (Blackwell experiment). We setup the sequential and present the equivalence of the sequential decision problem with a Markovian stochastic dynamic programming problem where the state is the decision maker's belief. Such recursive representation was firstly demonstrated by Hinderer (1970), then formalized in Easley and Kiefer (1988); Aghion et al. (1991) under slightly weaker assumptions.[18] We argue that the so derived value function has almost no use for describing the characteristics of a belief as it does not separate optimality and informational content of the same, thus yielding a problematic interpretation of the optimal policy. We define the dynamic Bayes risk as an admissible uncertainty function in the sense of DeGroot (1962) and show that the implied information function is traded-off one for one with expected utility at each belief state to determine the optimal policy.[19]

2.2.1 Primitives

Control

From a control perspective, an action is an Anscombe and Aumann act, that is a map from states Ω to lotteries over material outcomes $y \in Y$, Y being another compact subset of a metric space. The map of action and states into lotteries $\mathcal{L} : \mathcal{A} \times \Omega \to \Delta(Y)$ is continuous in both arguments.[20] Once an action is taken, nature (who knows the state ω) draws an outcome y according to $\mathcal{L}_{a,\omega} \in \Delta(Y)$. The Bernoulli utility index $\widetilde{u} : Y \to \mathbb{R}$ representing the agent's ranking over such material outcomes is taken as a primitive of the model. An action induces a map from states to (lottery) utilities, given by the section

$$u_a : \Omega \longmapsto \mathbb{R}$$

$$u_a(\omega) = \int \widetilde{u}(y)\,\mathrm{d}\mathcal{L}_{a,\omega}(y)$$

[18] Application of recursive techniques to bandit problems made the solution easy to characterize each period policy in terms of the Gittins index (Gittins and Jones (1974); Gittins (1979)). As the research in stochastic control discovered more and more advanced algorithms to solve bandit problems, see e.g. Bubeck and Cesa-Bianchi (2012); Auer et al. (2002), extensions of the original Robbins formulation were discussed mainly involving infinite and correlated arms.

[19] At a particular belief state, both expected utility and information are functionals over the action space. The former depends on its property as an AA act, the latter on its property as a Blackwell experiment.

[20] As Y is compact, $\Delta(Y)$ is metrizable As usual, the space of lotteries is endowed with the weak*-topology.

Continuity of the lottery mapping and of the integral function implies that u is (Borel) measurable and continuous in ω, a. Utility is further assumed to be uniformly bounded. For any probability measure $p \in \Delta(\Omega)$, I use without fear of confusion the function $u : \Delta(\Omega) \times \mathcal{A} \to \mathbb{R}$ to denote the expected utility

$$u(a,p) = \int u_a(\omega)\,\mathrm{d}p(\omega) = \int_\Omega \int_Y \widetilde{u}(y)\,\mathrm{d}\mathcal{L}_{a,\omega}(y)\,\mathrm{d}p(\omega)$$

Estimation

From an estimation perspective, an action is a Blackwell experiment, that is a collection of state-dependent measures over a common signal realization space S.

$$\mathcal{M}_a = (\mu_{a,\omega})_{\omega\in\Omega}\,, \text{ with } \qquad \mu_{a,\omega} \in \Delta(S)$$

Once action a is chosen nature draws a realization s according to measure $\mu_{a,\omega}$ and communicates it to the decision maker. This distribution being parametrized by ω provides the necessary signal-state correlation which permits learning to occur. It is assumed that $\mu_{a,\omega}$ is jointly continuous in action and states, where continuity is taken in the topology of weak convergence, that is for each function $h : S \to \mathbb{R}$ continuous, $(a_n, \omega_n) \to (a, \omega)$ implies

$$\int h(s)\,\mathrm{d}\mu_{a_n,\omega_n}(s) \to \int h(s)\,\mathrm{d}\mu_{a,\omega}(s)$$

2.2.2 Sequential and Recursive Solutions

The state space describing all the uncertainty that has to resolve at period 0 is given by $\Sigma = \Omega \times Z^\infty$, where $(Z, \mathcal{Z}, \nu)$ is a probability space and a realization $z_t \in Z$ in an information shock. At each period, independently of ω (and across periods) an information shock is drawn from ν. It then determines the realization of the signal for all possible state and action pairs through a function $s : \mathcal{A} \times \Omega \times Z$ such that $s_t = s(a_t, \omega, z_t)$.[21] Continuity of the signal maps is equivalent to having s continuous. In case $\mu_{a,\omega}$ all admitted a distribution function on $\mathbb{R} \supset S$, one can take $Z = [0,1]\,, \lambda$ the Lebesgue measure and

$$s(a,\omega,z) = \inf\{s \in S : z \leq F_{a,\omega}(s)\}$$

[21] This representation of objective uncertainty is taken from Aghion et al. (1991). It is convenient as it gives a common probability space that resolves the uncertainty in signal realization conditional on each state and action pair.

Σ is naturally endowed with the product $\sigma-$algebra, and the product measure $p_0 \otimes \nu^\infty$, where p_0 is the prior. Notice information shocks are by construction uncorrelated with ω, all correlation between s_t and ω being driven by the function s.

The sequential problem has the agent choose a plan $\widetilde{a} = (\widetilde{a}_0, \widetilde{a}_1, \ldots, \widetilde{a}_t, \ldots) \in L(\Sigma, \mathcal{A}^\infty)$ where $L(\Sigma, \mathcal{A}^\infty)$ denotes the set of random variables with state space Σ and sample space infinite action paths. A progressive measurability constraint comes in the form of requiring, for each $t \in \mathbb{N}$, $\widetilde{a}_t \in \sigma(Z^{t-1})$, where inclusion sign means measurability with respect to the RHS $\sigma-$algebra and, as usual $x^t = (x_s)_{s=0}^t$ and $\sigma(z^0) = \{\emptyset, \Sigma\}$ so that $\widetilde{a}_0$ is chosen not to depend on the state. Future actions can instead be made contingent on the information shocks that realized before period t. As s_t is a function of z_t, we can equivalently write a plan as a set of functions[22] $(\widetilde{a}_t : S^{t-1} \to \mathcal{A})_{t \geq 0}$, which is a more intuitive specification as it has plans depend on observables. Having $\widetilde{a}$ be a random variable that is correlated with ω (seen itself as a random variable) is convenient as it allows to write immediately the utility flow from a plan $\widetilde{a}$.[23] Let indeed

$$U(\widetilde{a}, p) = \mathbb{E}_{p \times \nu^\infty} \sum_{t=0}^{\infty} \beta^t u(\widetilde{a}_t, \omega) \tag{2.1}$$

We arenow ready to define the decision problem under study.

Definition 2.1. A **Multiarmed Bandit Problem (MAB)** Γ is a tuple $\Gamma = \{\Omega, \mathcal{A}, \beta, p_0\}$. The agent chooses a plan to solve

$$\max_{\widetilde{a} \in L(\Sigma, \mathcal{A}^\infty)} U(\widetilde{a}, p_0)$$

Let $v^*(p)$ be the value of this problem, namely

$$v^*(p) = \sup_{\widetilde{a} \in L(\Sigma, \mathcal{A}^\infty)} U(\widetilde{a}, p) \tag{2.2}$$

[22] I focus on action plans in pure strategies. Allowing for mixed strategies implies letting a_t being stochastic even conditioning on the signal realization. This will only add notational complexity, as we know that in a decision problem randomization is never optimal.

[23] The measure $\mathbb{E}_{p \times \nu^\infty}$ used to define (2.1) takes expectation of the correlated random variables $\widetilde{a}, \omega$.

Recursive representation

The Markov kernel, mapping prior into distirbutions over poserior induced by a given experiment, is a fundamental object in achieving the recursive representation of the MAB problem. We obtain it as follows. Let $\pi : \Delta(\Omega) \times \mathcal{A} \times S \to \Delta(\Omega)$ be the Bayes posterior mapping, namely,

$$\pi(p, a, s)(O) = \frac{\int_O \mu_{a,\omega}(\mathrm{d}s)\,\mathrm{d}p(\omega)}{\int_\Omega \mu_{a,\omega}(\mathrm{d}s)\,\mathrm{d}p(\omega)} \tag{2.3}$$

for any O in the Borel-sigma algebra on $\Delta(\Omega)$.[24] The Markov kernel μ_a has, for each $\mathcal{O} \in \mathcal{B}(\Delta(\Delta(\Omega)))$

$$\mu_a(p)(\mathcal{O}) = \int_\Omega \left(\int_{\{s:\pi(p,a,s)\in\mathcal{O}\}} \mu_{a,\omega}(\mathrm{d}s) \right) \mathrm{d}p(\omega) \tag{2.4}$$

We can now state

Proposition 2.1. *(Easley and Kiefer (1988), Aghion et al. (1991)) The value function v^* defined in 2.2 coincides with v, the solution to the functional equation*

$$v(p) = \sup_{a\in\mathcal{A}} \left[u(p,a) + \beta \mathbb{E}_{\mu_{p,a}} v(p') \right] \tag{2.5}$$

where $\mu_{p,a}$ is defined in (2.4). Further, v is continuous and convex.

The policy correspondence $\Phi : \Delta(\Omega) \to \mathcal{A}$ is given by

$$\Phi(p) = \left\{ a \in \mathcal{A} : \; u(p,a) + \beta \mathbb{E}_{\mu_{p,a}} v(p') = v(p) \right\} \tag{2.6}$$

specifies how the agent plays as a function of the currently held belief, which is a sufficient statistic for all information acquired along the experimentation path.

Belief as a state

Thinking about the solution of the sequential problem in its recursive representation (2.5) gives substantial simplification on how we interpret the attitude of the agent

[24] The space of measures over a compact metric space is metrizable, hence the object is well definite.

towards uncertainty.[25] However, having beliefs as the endogenous states presents some challenges. Beliefs are not capital or any state variable inside classic economic problems. They are not fixed, in the sense that there is nothing preventing the agent from changing his belief at the beginning of each period.[26] Why does not a player give himself the "best prior", hence choose $p^* \in \arg\max_p v(p)$, as the latter object is defined in the prior free decision problem Γ? The obvious answer is that the value $v(p)$ is not the value of "having belief p", but it is the value -the expected flow utility following an optimal path- of *p being the most accurate description of reality available* at the particular time: if one decides to be certain to be a multi-millionaire, he will still end up bankrupt. For this reason it is legitimate to call the state "information", rather than "belief", even though it creates confusion with the information function studied below.

The value of a belief does not distinguish between the "decisional quality" of states in its support (the utility they would deliver should them be true), and the "informational quality" of the belief himself (how much "uncertainty" it contains). To help visualizing this point, one can rewrite the recursive problem (2.5) by making the rational updating constraint explicit as

$$\begin{array}{cc} \max_{a \in \mathcal{A}} & u(p,a) + \beta v(p') \\ s.t. & p' \sim \mu_{a,p} \end{array}$$

Action parametrizes distribution of posterior, but *it is always the belief the element of the domain in the continuation value*: it is therefore not wrong to say that an agent tries to make herself believe that she is multi-millionaire while respecting the rational updating constraint.

The sequential representation derived above is an useful tool for solving numerically the problem. However the resulting value function has almost no use for describing the characteristics of a belief as it does not separate optimality and informational content of the same, thus yielding a problematic interpretation of the optimal policy. In the next subsection we use insights in DeGroot (1962) to derive the

[25] The complex structure of objective uncertainty coming in the form of a distribution over infinite sequences of signals as generated by a plan which in turns determines distribution over action paths and hence utility is simplified by just looking at a much simpler object, namely the Markov kernel $\mu_{p,a}$ induced by each action over posteriors in the immediately next period.

[26] This is particularly critical when the agent has to determine his prior, especially when such prior results from a structured conjecturing effort. The recursive representation has indeed the prior lose the special standing it had in the sequential representation, where we computed state-dependent measures of action paths conditional on each state, and then aggregate based on the prior.

relevant uncertainty function, describe why it conforms with the heuristic properties discussed in the introduction and how optimal policy trades off current optimality and information acquisition.

2.2.3 Uncertainty Functions

We follow the approach in DeGroot (1962) in defining information through uncertainty functions. The premise is heuristic: given a state space Ω there is a natural sense in which a belief $p \in \Delta(\Omega)$ describes a certain amount of uncertainty on the true value ω. What an uncertainty function attempts to make is to give a numerical quantification this "natural sense". The concepts of information and uncertainty are closely linked: whatever reduces uncertainty is information. The following definition makes this intuition formal.

Definition 2.2. An **uncertainty function** is a non-negative map $L : \Delta(\Omega) \to \mathbb{R}$. Given an uncertainty function L, the **information function** $I^L : \Delta(\Omega) \times \mathcal{A} \to \mathbb{R}$ is defined as

$$\begin{aligned} I_p^L(a) &= L(p) - \mathbb{E}_{p,a}(L(p')) \\ &= L(p) - \int_{\Delta(\Omega)} L(p') \,\mathrm{d}\mu_{p,a}(p') \end{aligned}$$

Where $\mu_{p,a} \in \Delta(\Delta(\Omega))$ is the Markov kernel (2.4) of experiment a evaluated at prior p. The information of an experiment $a \in \mathcal{A}$ evaluated at prior p is the difference between the uncertainty of the prior distribution p and the *expected uncertainty after having observed the realization of a.*

A minimal requirement for an uncertainty function to be admissible is that each experiment contains non-negative information. Clearly this does not imply that for all possible realizations of the experiment the uncertainty will decline (think of a case in which the Markov kernel is full support, namely each posterior is achieved with positive probability). The property refers more to a qualitative feature that an information function must have: "It is commonly felt, and often stated that an experiment can, at worst, contain no information about the problem at hand."

Proposition 2.2. ***(DeGroot (1962))*** *An uncertainty function is admissible if and only if it is (weakly) concave.*[27]

[27]The convex combination between two measures is defined as usual

$$(\alpha\mu + (1-\alpha)\mu')(A) = \alpha\mu(A) + (1-\alpha)\mu'(A)$$

for $A \in \Sigma$.

2.2.4 Solution through uncertainty functions

We now define a finction

Definition 2.3. Fix a prior-free sequential decision problem Γ. For each $p \in \Delta(\Omega)$, define

$$L^{\Gamma}(p) = \beta\left[\mathbb{E}_p\left[v\left(\delta_\omega\right)\right] - v(p)\right]$$

where v is defined in (2.2) and δ is the Dirac measure. $L^{\Gamma}(\cdot)$ is the **dynamic Bayes risk** associated to problem Γ.

The normalizing factor β is taken only for future expositional convenience. The interpretation is that the uncertainty of the decision problem at a particular ω is the loss in utility that occurs when the true state is ω and, rather than knowing it, the agent holds a belief p and hence acts (dynamically) optimally under such belief. The dynamic Bayes risk then aggregates all the losses received in each possible state under the distribution p. It can be interpreted as the expected flow utility loss which results from acting in an uncertain environment rather than in an uncertainty free environment *if the state ω was really drawn from p.* It also has a welfare interpretation as it is the amount (in utils) that an agent with information state p would be willing to pay[28] to operate in the same decision problem purified of subjective uncertainty.

The following proposition establishes that the dynamic Bayes risk is an admissible uncertainty function, provides its recursive representation that will be used in the applications of the next part to compute this object in simple cases that are numerically tractable, and establishes the key result on linear utility-information tradeoff. In the statement and proof (Appendix) it is implicit that we fix a prior-free sequential problem Γ.

Theorem 2.1. *i) L is an admissible uncertainty function and coincides with the unique solution of the functional equation*

$$L(p) = \min_{a \in \mathcal{A}} \mathbb{E}_p\left[u^*(\omega) - u(a,\omega)\right] + \beta \mathbb{E}_{a,p}\left[L\left(p'\right)\right] \tag{2.7}$$

ii) Let Φ be the policy correspondence defined in (2.6), $I^L : \Delta(\Omega) \times \mathcal{A} \to \mathbb{R}_+$ the information function associated to the dynamic Bayes risk and I^L_p its section at p. It

[28] To justify the β, the cost is paid at a period before the statistical game starts.

holds

$$a \in \Phi(p) \iff a \in \arg\max_{a\in\mathcal{A}} \left[u_p(a) + I_p^L(a)\right] \tag{2.8}$$

The characterization of the policy function (2.8) both concludes and constitutes the main result of this Section. Its conceptual contribution can be summarized as follows. The expected reduction in flow utility of an optimal experimentation plan compared to a counterfactual uncertainty-free world is a legitimate uncertainty function in the sense of DeGroot. It discriminates beliefs along dimensions that are reasonable foundations of the concept of uncertainty, such as utilitarian differences across belief-induced state realizations and the practical facility to obtain more accurate description of reality. A manipulation of the recursive representation of the repeated decision problem gives the decisional implications of the uncertainty function, as agents trade-off one for one current utility with uncertainty reduction. The next section shows how one can apply

2.3 Vaccination Example

In this Section we offer a numerical solution to Example 2.1. We have a twofold motivation: firstly we want to show how one can derive the uncertainty function (2.7) over a non-tivial (binary) state space and how to use it to compute the policy in the implied MAB problem. Secondly the model presented hereafter will be used as a motivating example for the next part. As for the first point, the state space will be continuous $\Omega = [0,1]$ though, to keep the model computationally tractable we will need to force information updating inside a conjugate model (the beta-binomial model), so that along any experimentation path the belief state will be identified with a pair of parameters. As will be clear from the discussion in later sections, the simplicity that makes such example computationally tractable comes at the cost of having most of the interesting channels of incentivation into social program problems being shut down. In subsequent sections we will indeed use this example as an instance in which the more complicated objects we introduce take a simple form.

2.3.1 Social utility and control optimality

The environment is described in Example 2.1. Period utility is assumed separable in distribution over health statuses and monetary revenues. Along the first dimension, the planner gets utility 1 for each healthy citizen and 0 for each ill citizen. He has to pay the cost of the plan which equals the fraction vaccinated times the per-unit net loss $q^* - q$. Parameter χ describes the relative importance of monetary vs. health outcomes, $\chi \to 0$ means the planner only cares about health of the population. Let $\mathbf{y}$ be the random variable describing the number of vaccinated agents, $\mathbf{i}$ the number of ill across non vaccinated. The state-action static utility is given by

$$\begin{aligned} u(\omega, q) &= \tfrac{1}{N}\mathbb{E}_{Bin(e^{-\gamma q},N)}\left[\mathbb{E}_{Bin(\omega,N-\mathbf{y})}\left(\mathbf{y} + (N - \mathbf{y} - \mathbf{i})\right)\right] - \mathbb{E}_{Bin(e^{-\gamma q},N)}\left[\chi \mathbf{y}\left(q^* - q\right)\right] \\ &= \tfrac{1}{N}\left[N - \mathbb{E}_{Bin(e^{-\gamma q},N)}\left[\mathbb{E}_{Bin(\omega,N-\mathbf{y})}\left(\mathbf{i}\right)\right] - \mathbb{E}_{Bin(e^{-\gamma q},N)}\left[\chi \mathbf{y}\left(q^* - q\right)\right]\right] \\ &= \tfrac{1}{N}\left[N - \mathbb{E}_{Bin(e^{-\gamma q},N)}\left[\omega\left(N - \mathbf{y}\right)\right] - \chi N e^{-\gamma q}\left(q^* - q\right)\right] \\ &= 1 - \omega\left(1 - e^{-\gamma q}\right) - \chi e^{-\gamma q}\left(q^* - q\right) \end{aligned}$$

which gives state contingent policies

$$\widehat{q}(\omega) = \max\left\{0, q^* + \frac{1}{\gamma} - \frac{\omega}{\chi}\right\} \tag{2.9}$$

It is assumed that $q^* + \frac{1}{\gamma} - \frac{1}{\chi} > 0$ so that we always have an interior solution. Notice $\frac{\partial \widehat{q}(\omega)}{\partial \omega} = -\frac{1}{\chi}$, optimal price strictly declines with ω: as the illness is more aggressive, the benefit of vaccinating is higher. As $\chi \to 0$, healthiness becomes more important (relative to monetary expenses) and so lower prices are charged. Substituting the optimal policy we get

$$u^*(\omega) = 1 - \omega + \chi e^{\frac{\omega}{\chi} - q^* - 1}$$

As the planner is risk neutral the *static* incentivation problem features certainty equivalence, that is for each $\mu \in \Delta(\Omega)$

$$\widehat{q}(\mu) = q^* + \frac{1}{\gamma} - \frac{\mathbb{E}_{\mu}[\omega]}{\chi}$$

gives the static best response function: whenever holding belief μ, a myopic planner would quote price $\widehat{q}(\mu)$.

2.3.2 Estimation

By risk neutrality the finite sample size is irrelevant for static optimality. It becomes relevant when we look at the estimation problem as finite larger qs will induce fewer people to vaccinate hence provide a larger size of the sample used to estimate ω.

Planner learns nothing from vaccinated citizens: they are healthy irrespectively of the state ω. Each non vaccinated citizen, however, becomes ill with probability ω

(independently one another). The planner makes therefore inference from the observation of a $Bin(\omega, N - \mathbf{y})$ random variable and updates his prior on the unknown parameter ω. Notice that the effective sample size $N - \mathbf{y}$ is itself the realization of a $Bin((1 - e^{-\gamma q}), N)$ random variable (parametrized by planners' action).

For each q the signal received is a pair of realized sample sizes and realized sick people. $S = [N] \times [N]$ is the space of signals that can be received and

$$\begin{array}{rcccc} s_1 \big| q & = & \#NoVacc & \sim & Bin\left((1 - e^{-\gamma q}), N\right) \\ s_2 \big| s_1, \omega & = & \#ill & \sim & Bin(\omega, s_1) \end{array}$$

It is clear that signal $s = (s_1, s_2)$ is a garbling over $s' = (s'_1, s'_2)$ if $s'_1 \geq s_1$: we obtain the same state-signal distribution by making the last $s'_1 - s_1$ health outcomes observations uninformative. Then there is a clear sense in which higher q is more informative, as it ranks more Blackwell informative signals according to first order stochastic dominance.[29]

For tractability we need to assume that prior $p_0 \in \Delta([0,1])$ belongs to the beta family: as the the binomial likelihood is conjugate with the beta distribution this means all posteriors (that is, information at each period) will be described by beta distributions.

As it is well known, after observing s_2 successes out of s_1 trials, a prior $B(\alpha, \beta)$ updates in

$$B(s_2 + \alpha, s_1 - s_2 + \beta)$$

It will be convenient to reparametrize the beta distribution defining $\nu = \alpha + \beta$ and $\mu = \frac{\alpha}{\nu}$: under this parametrization the scale parameter μ represents the mean of the distribution while ν is as a measure of the accuracy of estimation.[30] The posterior in the reparametrized version of the beta is

$$B\left(\frac{\mu\nu + s_2}{\nu + s_1}, \nu + s_1\right)$$

which reflects the fact that effective sample size always grows with time (we collect more and more observations), in a manner that does not depend on the particular timing of experimentation.

We now give the expressions for the Markov kernel $M_q : \mathbb{R}^2 \to \Delta(\mathbb{R}^2)$.

[29] If $p > p'$, $F_p(x) \leq F_{p'}(x)$ where F_p is the CDF of a binomial random variable with parameter p and any N.

[30] In the limit as $\nu \to \infty$, the measure approaches a Dirac on μ.

$$\nu'(\mu,\nu,q) \sim \nu + Bin\left(\left(1-e^{-\gamma q}\right), N\right) \tag{2.10}$$

that is the increase in the effective sample size ν is stochastic but does not depend on the unknown ω (hence does not provide information about it).[31] The posterior mean instead is going to have a distribution that depends on ω. In particular, given a realized sample size s_1, one has for $s_2 = 0, 1, \dots, s_1$,

$$\begin{aligned}
p_{\mu,\nu,s_1}(s_2) &= Prob\left(\mu' = \tfrac{\mu\nu+s_2}{\nu+s_1}\right) \\
&= \int_0^1 Bin_{\omega,s_1}(s_2)\,\mathrm{d}B_{\mu,\nu}(\omega) \\
&= \tfrac{\Gamma(\nu)}{\Gamma(\mu\nu)\Gamma(\nu-\mu\nu)} \int_0^1 \omega^{y+\mu\nu-1}(1-\omega)^{s_1-y+\nu-\mu\nu-1} \begin{pmatrix} s_1 \\ s_2 \end{pmatrix} \mathrm{d}\omega \\
&= \tfrac{\Gamma(\nu)}{\Gamma(\mu\nu)\Gamma(\nu-\mu\nu)} \begin{pmatrix} s_1 \\ s_2 \end{pmatrix} \int_0^1 \omega^{s_2+\mu\nu-1}(1-\omega)^{s_1-s_2+\nu-\mu\nu-1}\,\mathrm{d}\omega \\
&= \tfrac{\Gamma(\nu)}{\Gamma(\mu\nu)\Gamma(\nu-\mu\nu)} \begin{pmatrix} s_1 \\ s_2 \end{pmatrix} \tfrac{\Gamma(\mu\nu+s_2)\Gamma(s_1+\nu-s_2-\mu\nu)}{\Gamma(\nu+s_1)} \\
&= B(\nu-\mu\nu,\mu\nu)B(\mu\nu+s_2, s_1+\nu-s_2-\mu\nu) \begin{pmatrix} s_1 \\ s_2 \end{pmatrix}
\end{aligned} \tag{2.11}$$

Where Γ denotes the gamma function and $\begin{pmatrix} \cdot \\ \cdot \end{pmatrix}$ is the binomial coefficient. Expressions (2.10)-(2.11) give the expressions for the Markov kernel. Notice that the Markov kernel is continuous in actions and states: all measures have well definite and continuous densities for all positive q, we only need to check the limit as $q \to 0$, but as $q \to 0$, $\mu_q(\mu,\nu) \to (\delta_{\mu,\nu})$ in the topology of weak convergence.

Given this observation it is then immediate to check that the problem satisfies all the assumptions required in part 2: state space is compact, current utility utility is bounded and continuous and Markov kernel is continuous.[32]

The planner then chooses a progressively measurable stochastic process for prices to solve

$$\max_{q \in L(\Sigma,\mathbb{R}^\infty)} \mathbb{E}_{p\otimes\lambda^\infty} \sum_{t=0}^{\infty} \beta^t u(\omega, q_t)$$

[31] We can immediately see an extension would have γ unknown as well so that also the observation of s_1 is informative about the underling reduced form model. In the language of part 4, the model is degenerate along the statistic sufficient for treatment, only outcome types are unknown.

[32] It remains to show that it is without loss of generality to restrict the set of prices to a compact interval.

where λ is the measure over information shocks that determine the price-state conditional signal realization according to (2.10)-(2.11).

Numerical solution of the uncertainty function

By Theorem (2.1) we can write the dynamic Bayes risk for this problem as the solution to the following functional equation

$$L(\mu,\nu) = \min_{q\in Q} \mathbb{E}_{\mu,\nu}\left[u^*(\omega) - u(q,\omega)\right] + \beta\mathbb{E}_{M_q(\mu,\nu)}\left[L(\mu',\nu')\right]$$

where, for $Q=\mathbb{R}$, period loss is[33]

$$\mathbb{E}_{\mu,\nu}\left[u^*(\omega) - u(q,\omega)\right] =$$

$$= \mathbb{E}_{\mu,\nu}\left[\chi e^{\frac{\omega}{\chi}-q^*-\frac{1}{\gamma}} - e^{-\gamma q}\omega + \chi e^{-\gamma q}(q^*-q)\right] =$$

$$= \chi e^{-\gamma q^*-1}\mathbb{E}_{\mu,\nu}\left[e^{\frac{\gamma}{\chi}\omega}\right] + e^{-\gamma q}\left[\chi(q^*-q) - \mu\right]$$

and $\mathbb{E}_{M_q(\mu,\nu)}$ denotes expectations with respect to the measure described by (2.10)-(2.11) derived above, namely

$$\mathbb{E}_{M_q(\mu,\nu)}\left[L(\mu',\nu')\right] = \sum_{n=0}^{N} Bin_{(1-e^{-\gamma q}),N}(n)\sum_{y=0}^{n} p_{\mu,\nu,n}(y)\, L\left(\frac{\mu\nu+y}{\nu+n},\nu+n\right)$$

with

$$p_{\mu,\nu,n}(y) = B(\nu-\mu\nu,\mu\nu)B(\mu\nu+y, n+\nu-y-\mu\nu)\begin{pmatrix} n \\ y \end{pmatrix}$$

We solve numerically the model using the following values for the known parameters: $\gamma = 0.2$ $\chi = 0.3$, the discount factor β can take the three values 0.1, 0.5 and 0.95 reflecting different importance of the future periods. We also consider two different action sets among which prices can be chosen: either prices belong to a discrete set ($Q = \{0, L, H\}$ corresponding to free vaccine, low and high price), or $Q = \mathbb{R}_+$, any positive price can be quoted.

[33] $E_{\mu,\nu}\left[e^{\frac{\gamma}{\chi}\omega}\right]$ has an analytic expression being the moment generating function of a beta μ,ν evaluated at $\frac{\gamma}{\chi}$. It is of no particular interest.

Figures 2.1 and 2.2 plot the uncertainty and policy functions for different values of β in case the incentive set contains 3 prices or a continuum, respectively. Uncertainty is on the LHS of each panel.[34]

Policies are plotted on the RHS. On the horizontal axis is the value of the mean parameter μ, while on the vertical is the precision parameter ν on a log scale. Different colors represent different optimal price associated to a (μ, ν) pair, with darker colors (towards blue) represent lower prices. As $\nu \to \infty$ we have $B(\mu, \nu) \to \delta_\omega$ in the topology of weak convergence, therefore the upper stripe ($\log(\nu) \geq 4$) in each figure represent[35] the static policy when $\omega = \mu$ is known, (for continuous action space, the policy is given by (2.9)). We see that as agent becomes less certain of μ, that is decreasing ν for fixed μ, optimal price increases. This violates the certainty equivalence property of the decision problem and reflects the wedge introduced by information acquisition, since quoting higher price increase the (expected) size of sample that provides information about ω. At lower ν (that is, where agent is not very confident about μ), he will induce oversampling by setting higher prices. Distortions are more pronounced when the discount factor is higher (notice for $\beta = 0.1$ the policy in the $3-$action case almost obeys certainty equivalence) as the benefit of information is discounted more.

Uncertainty does not decrease as we enlarge the action space. This is because an expansion of the action set does not simply increase the possibilities to learn but also complicates the correlation between states and static policy: with 3 actions we only need to know a $3-$atoms partition of the state space, while a continuous action set makes the optimal action a function of the true state. Therefore increased possibilities to learn come together with a more complicated domain of relevant uncertainty, making the two environment not unambiguously comparable.

Uncertainty is instead uniformly decreasing with N: current policy and utility are unaffected by risk neutrality, but higher N make the realized sample larger.[36] This is equivalent to a larger μ in the previous example, it makes every action more informative hence the whole environment less uncertain. It should be noticed that, as $N \to \infty$, $\overline{y} \to \omega\left(1 - e^{-\gamma q}\right)$ in probability by the law of large numbers. As the RHS is an invertible function of ω for each $q > 0$ this means that for every positive price

[34] Concavity of the surfaces does not follow immediately from concavity of the uncertainty, as the latter holds in the space of probability distributions and convex combinations of beta are not betas with convex combinations of parameters.

[35] Up to finite grid approximation.

[36] We can simply through away all observations related to citizens labeled from N to $N' > N$, have the same decision problem but Blackwell inferior experiment.

quoted an infinite sample makes the observe success rate reveal ω. Therefore $I_p(q) = \mathbf{1}_{q>0} U(p)$, all actions are equally informative hence the representation in Proposition 2.1-*iii*) implies a trivial dynamics for the optimal solution: at the initial period we set $q = q^*(\overline{\omega}_p) > 0$ static optimum,[37] and from the second period onward the state is revealed and we act according to (2.9) under complete information. Triviality of the exploration-exploitation trade-off under identification rather than estimation concerns is further discussed in subsection 2.4.5.

[37] If $q^*(\overline{\omega}_p) = 0$ no solution existed as the "optimal" price is to set $q = \epsilon > 0$ as small as possible.

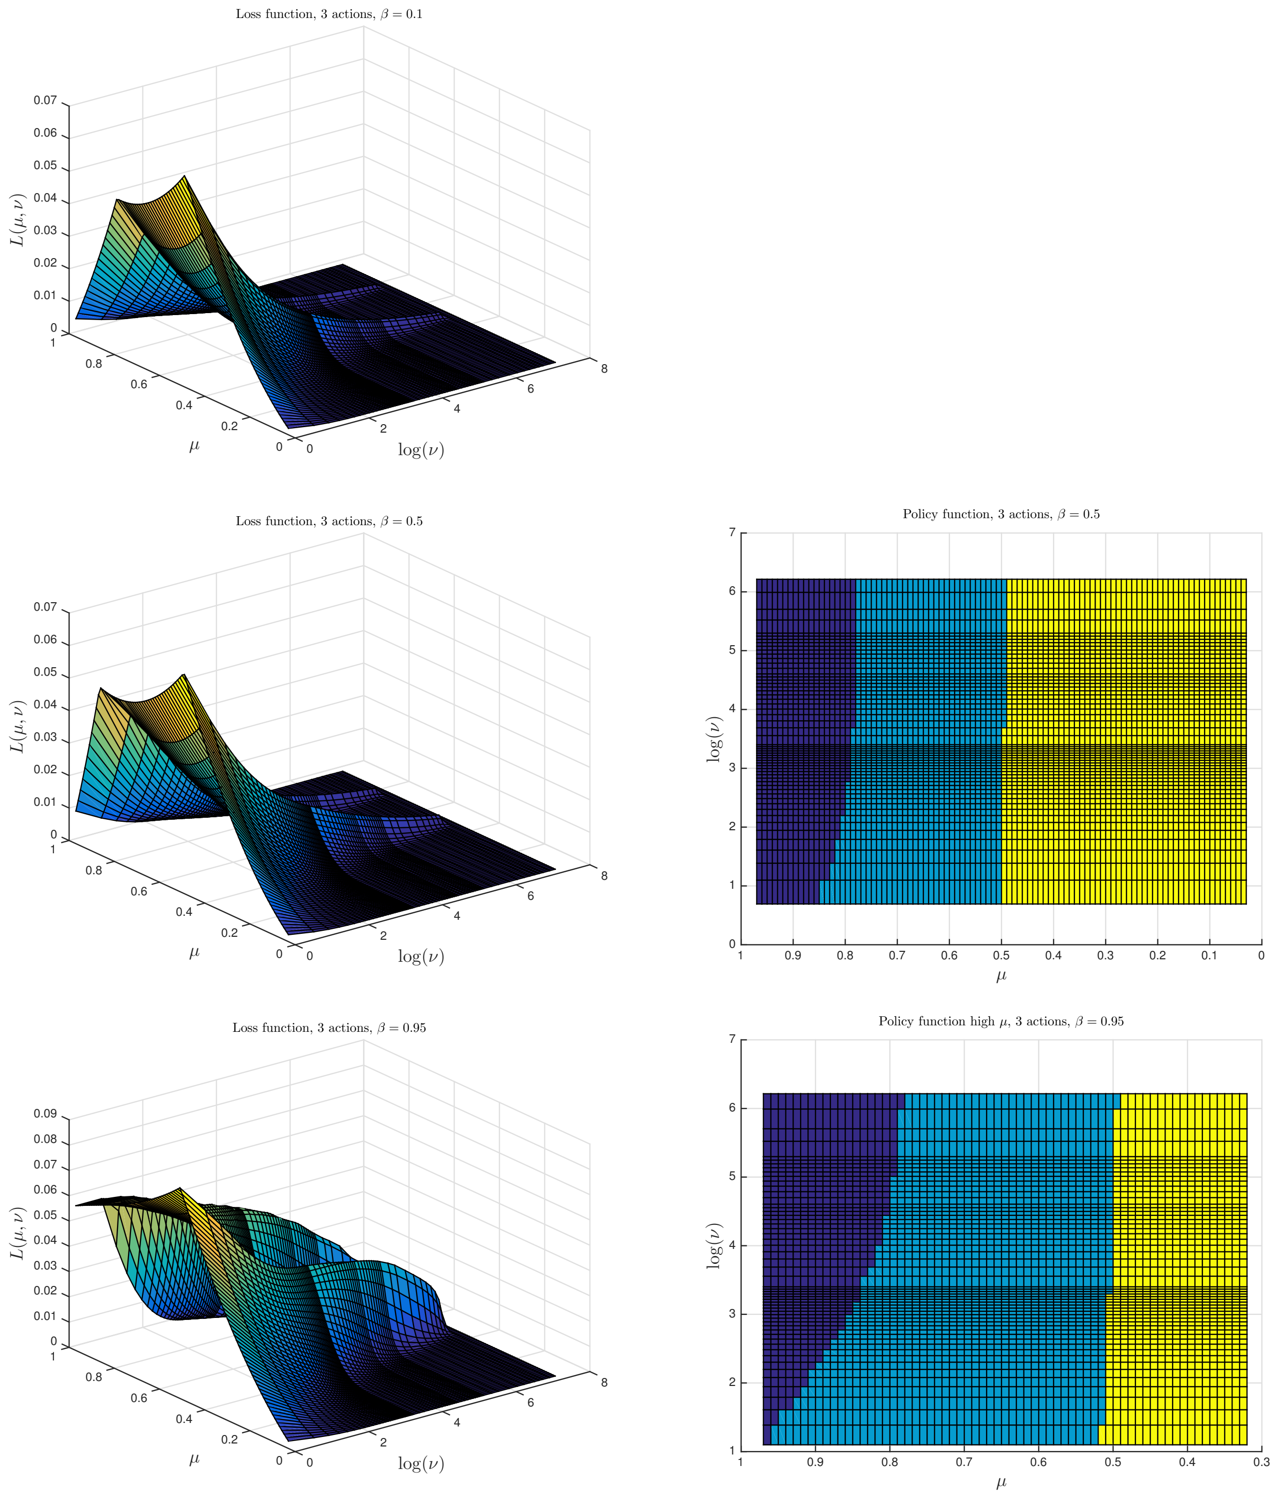

Figure 2.1: Loss and policy functions for 3 actions (different β)

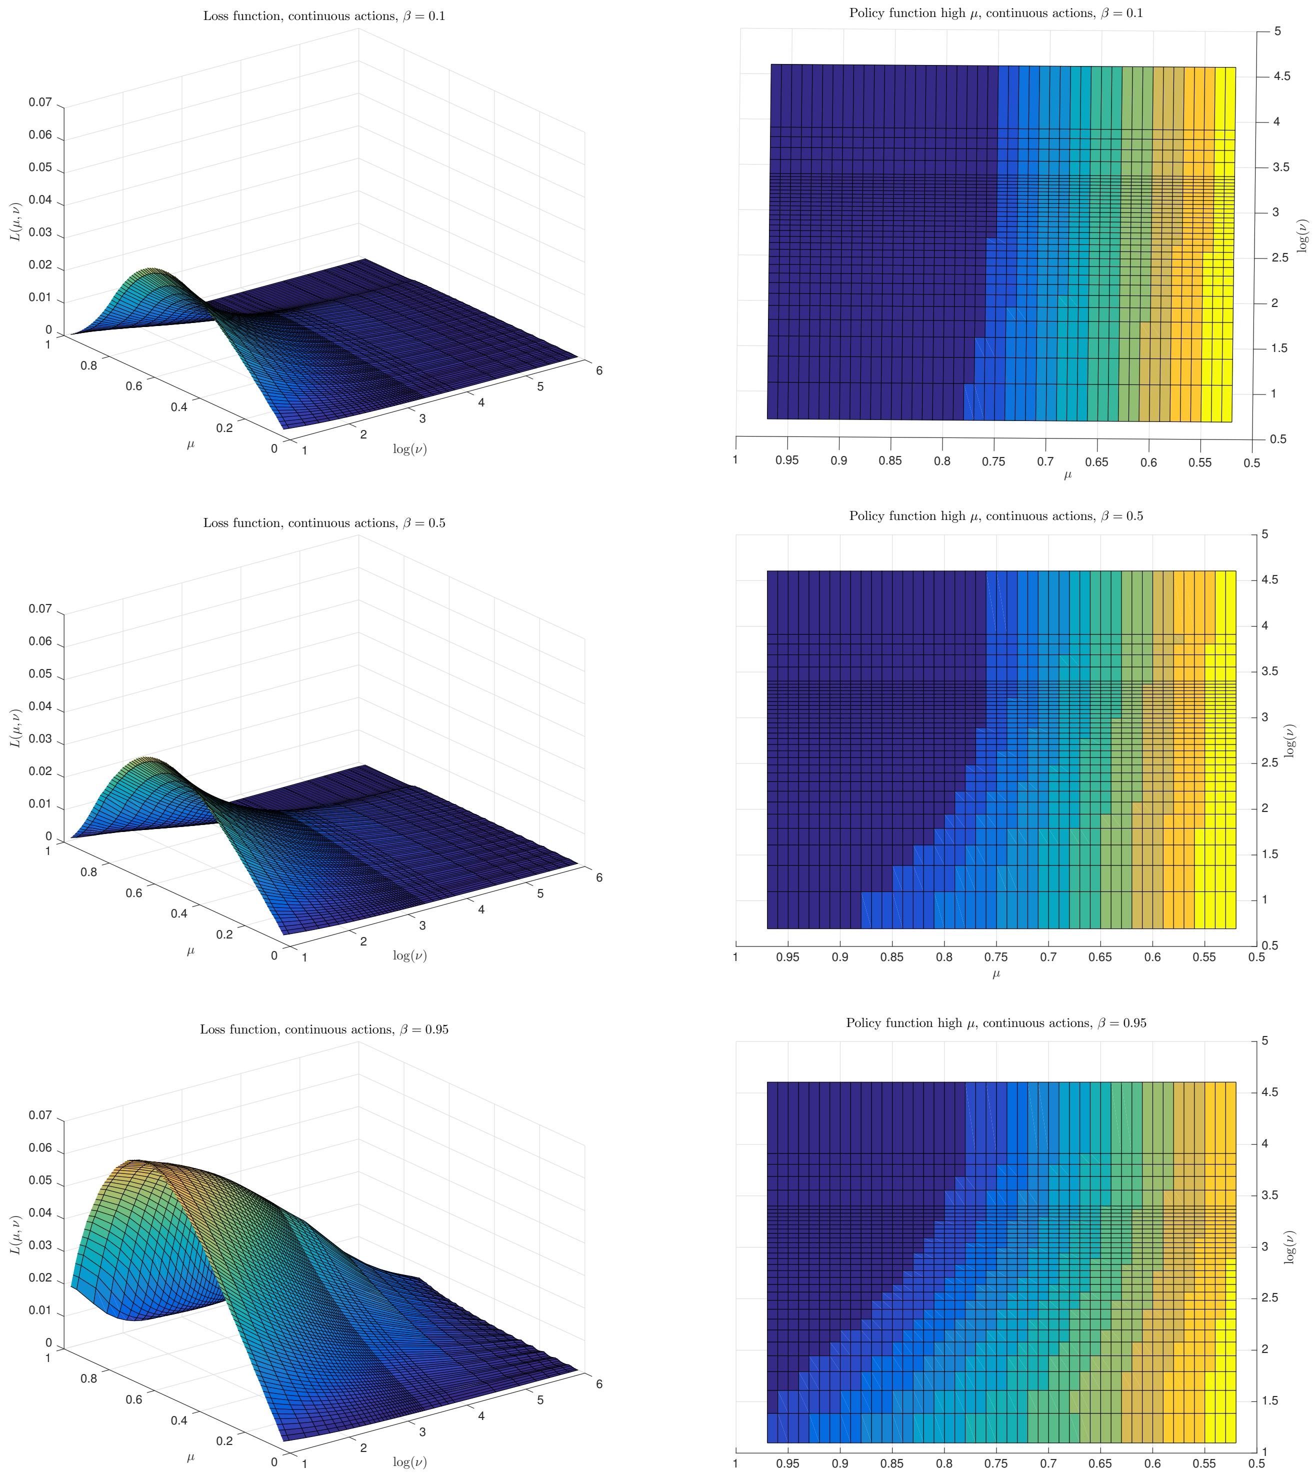

Figure 2.2: Loss and policy functions (for high μ) for continuum of actions

2.4 The general model: State and Action Space

The remainder of this Chapter is devoted to setting up problems of incentivation into a social program as a particular class of the MAB problems introduced in Section 2; the vaccination example solved in Section 3 will be a particular tractable case in which most of the channels that make the problem interesting are shut down for the sake of tractability. To use a language that conforms with the literature on estimation of treatment effects we will be as adherent as possible to the terminology used in Imbens and Rubin (2015) in terms of statistical assumptions (SUTVA, unconfoundedness, etc.), and objects of interest (superpopulation, propensity score). This Section carefully reviews the modelling assumptions that are kept implicit from now on, then specifices the bandit problem by deriving the state and action spaces that are relevant for the application.

We have a "circularity" problem when presenting state and action space: loosely speaking, the state is the distribution of individuals' responses to each possible stimulus in the action set (in terms of treatment choice and messages sent), therefore the state space is only definite with respect to an action set, which in turns incentivizes along dimensions that are encoded by the individual state. We solve this circularity by starting to define the (social) state space holding a generic action space fixed (Section 2.4.2), then fill that gap (Section 2.4.3). We offer weak sufficient conditions for superpopulation implementation of the class of BDM mechanisms; individual response to all such mechanisms are summarized by a single-dimensional empirical object (the reservation price). This is extremely convenient, however we also discuss potential failures driven by the complexity of the object being auctioneed (a treatment lottery). In Section 2.4.4 we introduce the concept of selection *regimes*, class of selection mechanisms that are ranked according to an intensity index (higher intensity selections have each type participate to treatment with higher probability), and rich enough find for each nondegenerate model and propensity score level an elements in the class that associates them. Section 2.4.5 discusses the informational content of BDM selection mechanisms, making the point that the selection properties of a sampling intensity are not necessarily related to its estimation power, and in particular that guaranteeing truthful revelation does not necessarily imply that planner gets to known the reservation price.

2.4.1 Modeling Assumptions

We now proceed to list and discuss the 5 critical assumptions that are maintained throughout this paper, from the most to the least restrictive.

1. Stable environment: the action set in invariant through time and the population alive each period is an $N-$dimensional random draw from an immutable super-population.

2. Outcome conditional on no treatment is known and, if stochastic (that is not constant across the no-treated population), does not depend on individual characteristics that determine treatment selection or outcome conditional on treatment.

3. All strategic interaction channels are shut down, both in the treatment selection and outcome phases.

4. There are no observable covariates, hence measurability of the stimulus implies that everyone is subject to the same stimulus.

5. Treatment and outcome spaces are binary.

The first assumption is the most fundamental one that makes policymaking a particular type of bandit exploration. This assumption is particularly critical for environments in which individual decisions depend on beliefs as citizens are not allowed to learn. Further failures of stationarity would occur if planner was able to influence the distribution of characteristic through its action: think for instance of a disease-killing intervention or of promoting school to have a better educated population in the future which is also likely to be more sensible to the importance of educating their own children.

The second assumption gives the treatment status a qualitative feature that is absent in most of the standard models of treatment effect estimation where treatment and control have no intrinsic meaning and are simply two different statuses. Here, selecting into treatment provides information about the potential outcome type that is not revealed otherwise. Notice that under this definition in the vaccination example treatment corresponded to *not* being vaccinated, as uncertainty was related to that phase.

The third is the composition of two assumptions that are popular in the literature of estimation of treatment effect. No strategic interaction in the outcome phase is a (weak version of) the Stable Unit Treatment Value Assumption (SUTVA)[38] stating

[38] Actually, this is a weaker version compared to what given, for instance in Imbens and Rubin (2015). Here it is allowed that different agents use treatment with different intensities through their effort. This point is related to unconfoundedness assumption discussed in Section 2.6. Think of SUTVA in this weaker form as an assumption to exclude all sorts of treatment externalities.

that an individual's outcome depends exclusively on its treatment status: there are no "treatment externalities". No strategic interaction in the assignment phase is implied by an individualistic assignment assumption and will rule out competitive assignment mechanisms in which the treatment status depends on the identity of other citizens that are alive. A recent literature (Kremer et al. (2014); Mansour et al. (2015, 2016) Mansour et al. (2016)) focuses on agents' interaction inside a bandit problem through an arbitrary Bayesian game which allows to study a a three-way tradeoff between exploration, exploitation, and incentive provision to agents who are myopically interested in exploitation. The objective of the planner is to induce incentive compatible exploration, that is to make agents explore different arms[39] under the constraint that individuals will act for their own interest. The crucial assumption that makes those contributions unusable in my setting is outcome of pulling an arm is independent on the identity of the individual who pulls the arm: in the class of problems under study, however, it is crucial that the characteristic that determine selection also determine the outcome on the arm pulled (social program selected).

The fourth assumption is critical but is taken only to simplify the exposition. We can see this construction as a "building block" of more general model where the incentivation scheme needs to prescribe a citizen-specific stimulus assignment subject to measurability with respect to the observable partition. Observable covariates are of two types. Demographic observables on which stimulus assignment can be made contingent create a segmentation (in the sense of Bergemann et al. (2015)) in the market for treatment lotteries; covariates also partition the population into social groups (classes, villages) that provide natural experimentation units. Issues of external validity will require the planner to conjecture about the correlation of responses across atoms of the partition generated by the observable characteristics.

The fifth assumption is just for convenience, not to carry over a whole conditional distribution over outcomes.

2.4.2 The state space: speculation and reduced form model

There is a super-population of individuals, represented by the continuum $[0, 1]$, endowed with the Borel $\sigma-$algebra and the Lebesgue measure. Suppose for the moment the planner is given a set of actions $\mathcal{A}$, a compact subset of a metric space. $\mathcal{A}$ collects all possible stimuli that the she can conceive of enforcing to incentivize

[39] Though returns are in general allowed to be correlated.

individual decision to treat and the disclosure of private information. The individual type space is

$$T = ([0,1] \times M)^{\mathcal{A}} \times Y \tag{2.12}$$

Each individual $i \in [0,1]$ is associated to a type t_i, immutably through time. In what follows, I will refer to elements of the super-population as individuals and to elements of the population as citizens. An individual becomes a citizen (carrying with him its type) in periods where the sampling procedure makes him alive. If an individual is called to act in different periods he will give the same response to stimuli. Per (2.12), a type $t = (\tau_t, y_t) \in T$ specifies

- A treatment map $\tau : \mathcal{A} \to [0,1] \times M$, where the first coordinate gives as a function of stimulus a the probability with which type t selects into treatment, while the second coordinate gives the messages that are sent during either the treatment assignment phase or the post-treatment phase and enter the information set of the planner.

- A potential outcome $y \in Y$ describing how type t reacts to treatment, namely the health status that would result *should she be treated*. By assumption, non-treated get a deterministic health status so that the super-population type needs only to record the treatment-conditional one.[40] We take $Y = \{0,1\}$, bad and good health status (Assumption 5 above).

Notice that specification (2.12), altough pretty general, already embeds some of the standard assumptions taken in the literature of treatment effect estimation. The stimulus-conditional selection intensity (first coordinate of τ) does not depend on other citizens' identity and actions. This is an individualistic assignment assumption, and it will exclude competitive assignment mechanisms such as auctions. The potential outcome type y is assumed independent on i) how the individual ended up treated, namely, the stimulus she had responded to and ii) the treatment status of other citizens. The first assumption is an instance of policy invariance: an individual will not alter his post treatment behavior depending on whether he was, say, forced into treatment or had to freely choose it. Point ii) is instead a version of SUTVA.

[40]Thus, contrary to the treatment literature I will have the potential outcome to be a random variable, not a random vector.

What is not restricted by the specification (2.12) is the joint behavior of treatment type τ_t and outcome type y_t, which would induce in general a failure of the unconfoundedness assumption (see Section 2.6), and have the planner try to stimulate the "right citizens" into treatment.

As $\mathcal{A}$ and M are compact subsets of a metric space then so is T and we can define a $\sigma-$algebra over it, call it $\mathcal{T}$. The following object is then well definite

Definition 2.4. The **social state** $\omega^{SS} \in \Delta(T)$ is the push-forward probability measure induced on $\mathcal{T}$ by the type mapping, that is

$$\omega(A) = \lambda(\{i : t_i \in A\}) \ \forall A \in \mathcal{T}$$

The social state characterizes the super-population response in terms of selection intensity, health outcome and signals sent as a function of each possible stimulus. Given as it is, this object has little structure: it is defined by how it works, that is by how it maps feasible actions into super-population outcomes. This is of course convenient when it comes to choose an action conditional on knowing the model, but makes probabilistic conjecturing about the likely shape of such models a fairly difficult task. A theory on the joint behavior of determinants of decision and post-treatment behavior in the form of (structural or reduced form) model is needed to discipline the social state.

Definition 2.5. A **speculation structure** (X, f) is a set of individual characteristics X and a function $f : X \to T$ such that $t = f(x)$, characteristics determine the type. A **structural model** is a probability measure over speculation characteristics $\omega^{ST} \in \Delta(X)$.

Clearly one can take $X = T$ and have f be the identity and have a sort of "super-population model".[41] However the purpose of constructing a speculation structure is to reduce the dimensionality of the state space by giving a shape the individual pre and post treatment decisional environment.

We give now an example of a speculation structure; it is taken on purpose quite rich in terms of the number of characteristics considered, though it wouldn't be hard

[41] Notice Definition 2.4 does not use the term "model" which is reserved to outcomes of a conjecturing exercise over the space of social state.

to make it even more complicated. Most examples we will introduce later are slight modifications of this one.

Example 2.2. If individuals don't get treated, they get ill for sure (bad outcome). If individuals get the treatment, then they choose an effort inside an individual action space E by solving a subjective maximization problem. Good outcome results if and only if the chosen effort e^* exceeds an individual outcome type threshold $\overline{e}$. Agents believe that their likelihood to have successful treatment is $\Phi(e,\gamma) : E \times \Gamma \to [0,1]$, with Φ_1, Φ_2 strictly positive. They have a belief $p \in \Delta(\Gamma)$ and choose e to solve

$$\max_{e\in E} \int_{\Gamma} \Phi(\gamma, e) \cdot u_1 \mathrm{d}p(\gamma) - c(e) \tag{2.13}$$

u describes risk rankings (u_0 is normalized to 0 for all citizens, so that $\Phi(\gamma, e) \cdot u_1$ is the expected utility conditional on γ being true). To have a fully parametric model one can assume $\Phi(z,y) = (zy)^{\kappa}$, $c(z) = z^{\zeta}$, $p = p(\lambda)$. Individuals are moreover borrowing constrained and have a money endowment of $\overline{m}$. When considering whether to buy treatment at given price m, they compare the indirect utility of problem (2.13) with $g(\overline{m} - m)$, g being strictly increasing, say $g(z) = z^{\xi}$. If feasible stimuli are, as in 3.2, a set of prices at which agents can choose whether to buy treatment or not, the characteristic vector would be $x = [\overline{e}, \kappa, u_1, \zeta, \lambda, \xi]$. It describes outcome types, attitudes towards risk, disutility from effort, belief types, and utility from money respectively. It is immediate to see that this is indeed a speculation structure. Given x, one has for each possible action (price q) the following maps, assuming $M = \emptyset$,

$$\tau_x(q) = \mathbf{1}\left[\max_{e\in E}\left\{\int_{\Gamma} \Phi(\gamma, e) \cdot u_1 \mathrm{d}p(\gamma) - c(e)\right\} \geq (\overline{m} - q)^{\xi}\right]$$

$$y_x = \mathbf{1}\left[\arg\max_{e\in E}\left\{\int_{\Gamma} \Phi(\gamma, e) \cdot u_1 \mathrm{d}p(\gamma) - c(e)\right\} \geq \overline{e}\right]$$

A structural model is in this example a point $\omega^{ST} \in \Delta(\mathbb{R}^6)$. A specification gives mathematical structure to assumptions about individual's subjective decisional environment and the "true" environment that determines health outcomes. For instance, one can assume that although agents think they can alter their likelihood of success, everything is determined by other factors that the individual cannot control.[42] This is made by having $\overline{e}$ taking only two values a, b with $a < \min E < \max E < b$.

[42] What we implicitly did when setting up the vaccination example.

Alternatively, assumptions on private information by agents are made by specifying the correlation between outcome type $\overline{e}$ and beliefs λ.[43] Assumptions like "rich people are more lazy" is just $\overline{m}$ and ζ being positively correlated, and so many other qualitative statements about individual behavior are translated into the statistical properties of the structural model.

Essentially, the speculation structure corresponds to an economic theory of individual decision which disciplines citizen's behavior and allows to simplify (reduce dimension) the domain of uncertainty of the planner. Speculation structures include characteristics that may be neither identified in the data nor (taken individually) relevant for the outcome model.[44]

Definition 2.6. Given a speculation structure X, a statistic c is a function $c : X \to C$. A statistic is **sufficient for treatment** if the outcomes of the treatment phase depend through x only through its value $c(x)$, that is $c(x) = c(x') \implies \tau_x = \tau_{x'}$, the latter being a functional equality. It is **minimally sufficient for treatment** if it is a function of all statistics sufficient for treatment. Sufficiency (and minimal sufficiency) for messages and outcome are defined in the same way. Sufficiency without qualification means sufficiency for all phases.

For Example 2.2, immediate inspection of the treatment function gives

$$c_\tau(x) = \left[\max_{e\in E}\left\{\int_\Gamma \Phi(\gamma, e)\cdot u_1 \mathrm{d}p(\gamma) - c(e)\right\}, \overline{m}, \xi\right]$$

is sufficient for treatment; it will not be minimally sufficient since

$$c'_\tau(x) = c'(c_\tau(x)) = \min_q \left\{q \in \mathbb{R} : \max_{e\in E}\left\{\int_\Gamma \Phi(\gamma, e)\cdot u_1 \mathrm{d}p(\gamma) - c(e)\right\} \geq (\overline{m} - q)^\xi\right\}$$

[43] Example 2.9 elaborates on this and has agents observe a noisy signal of their outcome type $\overline{e}$ thus evaluating likelihood of success on based on the updated belief.

[44] The celebrated Marschak's Maxim (Marschak (1953)) states "All that may be required for policy analysis are combinations of subsets of the structural parameters, corresponding to the parameters required to forecast particular policy modifications, which are often much easier to identify."

will determine the response of an individual with characteristic x to any possible price stimulus as $\tau_x(q) = \mathbf{1}_{q \leq c'_\tau(x)}$. Also,

$$c_y(x) = y_x$$

is minimally sufficient for outcome. This is general: by the policy invariance assumption we do not need to take into account of the possible effects that feasible actions have on treatment-conditional outcome, therefore the minimal sufficient characteristic for outcome just coincides with the outcome function in a speculation structure. The function $\widehat{y} : X \times \mathcal{A} \to Y$ gives the final health status as a random variable for an agent with characteristic x subject to stimulus a.

$$\widehat{y}_x(a) = \begin{cases} y_x & w.p.\ \tau_x(a) \\ y_{NT} & w.p.\ 1 - \tau_x(a) \end{cases}$$

in case $\tau_x(a) \in \{0, 1\}$ this is deterministic.

Now we can finally define the state space of the social choice problem.

Definition 2.7. A **reduced form model** ω^{RF} is a joint distribution over minimally sufficient statistics. In particular, $\omega^{RF} \in \Delta(U \times Y \times S) = \Delta(\Omega^{RF})$ where $u : X \to U$, $s : X \to S$ are minimally sufficient for treatment and signals, respectively. Reduced form types are realizations of the random vector $\mathbf{u}, \mathbf{y}, \mathbf{s}$, defined on the common probability space $(X, \mathcal{X}, \omega^{ST})$. For $A \in \mathcal{U} \otimes \mathcal{S} \otimes \mathcal{Y} := \mathcal{B}(\Omega^{RF})$,

$$\omega^{RF}(A) = \omega^{ST}(\{x : [u(x), y(x), s(x)] \in A\})$$

It will be convenient to write

$$\omega^{RF}(\mathrm{d}u, 1, \mathrm{d}s) = \sigma^\omega(s, u)\, \omega_{S|U}(\mathrm{d}s|u)\, \omega_U(\mathrm{d}u) \tag{2.14}$$

Where $\sigma^\omega(s, u) = \mathbb{P}_{\omega^{RF}}(\mathbf{y} = 1 | \mathbf{s}, \mathbf{u} = s, u)$.

It is important to stress that a reduced form model does not only provide the information sufficient for static optimality (an outcome model $\omega^O \in \Delta(U \times Y)$). In a dynamic setting, the planner also needs a model to describe how signals that are (possibly) observed are jointly distributed with outcome relevant variables, as this correlation will be used in the estimation phase. By assumption the set of messages contains Y, and for some mechanisms (explicit BDMs), it will also include U; however we to keep those separated, at least conceptually.

The vaccination example of the previous section postulated directly a reduced form model. The sample space of minimal sufficient statistics was $C = \mathbb{R}_+ \times \{0, 1\}$ with

$$\tau_c(q) = \{\mathbf{1}_{q \leq c_1}\}, \; y_c = c_2$$

and

$$\omega^{RF}([0, x], \{0\}) = F_\gamma(x) \cdot \omega \tag{2.15}$$

where F_γ is the CDF of an exponential γ and ω was the unknown illness aggression parameter that characterizes uniquely the reduced form model. Notice that the first element in C must give selection probability as a *function* of possible stimuli (in this case, prices quoted). In this case the second coordinate is simply the outcome type.

The state space for planner problem is the set of reduced form model; I will denote its elements with ω (without RF superscript).[45] At period 0, before the first action is taken, the planner needs to form a conjecture $p \in \Delta\left(\Omega^{RF}\right)$. Models in the support of the conjecture may differ substantially: in the case of vaccination we just had uncertainty about a "average illness" but the structure is flexible enough to accommodate for this sort of parametric vs. model uncertainty; the requirement is only that the support of the prior conjecture (hence the set of perceived possible social states) is a compact subset of $\Delta\left(\Omega^{RF}\right)$.[46] Given such prior and the stationary nature of the problem, all posteriors will be derived from the Bayes rule based on the observations received, without further conjecturing effort being required.[47]

[45]The "true" object of uncertainty (the social state) does not put any discipline on response to different stimuli apart from that directly imposed by the planner's belief. The speculation structure and associated structural model are just used as conjecturing mediators, as they give a mathematical description of the subjective environment in which agents operate before and after choosing participation to a social program.

[46]We need to guarantee that $\Delta\left(\Omega^{RF}\right)$ is compact. In the relevant application of price incentivized mechanism with effort possibly observed, both U and S will be intervals in the real line.

[47]The approach is not immune to criticisms related to prior dependence of Bayesian estimation methods, especially if the prior conjecture is degenerate along certain dimensions of the reduced form model (as in the vaccination example). The planner however needs a conjecture on the relevant state to make a first period decision, it is not clear why he should disregard this conjecture in the estimation phase. Empirical Bayes methods (Robbins (1985)) may provide useful tools to overcome this criticism. Fessler and Kasy (2016) propose an estimation approach based on the empirical Bayes paradigm and apply it to models of labor demand.

So far we took the set of actions $\mathcal{A}$ as given. Relevant uncertainty is defined by the possible use the planner can make of it (selection policies) and by what is possibly observed (monitoring content of an action). The next subsection fills this gap.

2.4.3 The action space: non competitive price incentives

Per (2.12), an action has two dimensions: it prescribes a stimulus, which determines individual response into treatment, and describes the type of data that are received during the treatment assignment phase and the post treatment phase.

Each agent perceives treatment as a subjective environment in which he is (possibly) allowed to take decision which, jointly with some personal characteristics will determine participation-conditional health outcome. We focus on coercive stimuli, namely stimuli that operate by restricting the individual action set:[48] the planner holds a set of "treatment tickets" and allocates those through non-competitive pricing mechanisms.[49]

Subjective evaluation of the treatment environment is going to determine a reservation price, the maximal amount of money the want (or are able) to pay in order to change their status from "not treated for sure" to "treated for sure". Segmentation along the reservation price characteristic is easy to obtain under weak assumptions[50] on individual behavior through Becker-DeGroot-Marshak (BDM) mechanisms which were originally used to elicit the certainty equivalent of a monetary lottery and then proposed as a treatment assignment mechanism in CPS. The practical easiness to condition treatment on this unobserved characteristic comes at the cost of losing a clear theoretical connection between such characteristic and other determinants of success, especially if quasilinearity assumption is abandoned and one cannot interpret the reservation price as the "value from treatment".

Treatment lotteries

We will use subjective lotteries to describe the individual perception of the treatment environment. This introduces notation that I will keep also when describing the

[48] Extensions on non-coercive incentivation are discussed in subsection 2.7.1.

[49] Extensions on non-coercive incentivation are discussed in subsection 2.7.1.In the vaccination example this was the vaccine shot. So far it is not specified whether planner is subject to a sort of budget constraint in the form of a maximum number of treatments available, but Proposition 2.5 addresses this issue.

[50] The discussion after actually points that assumptions needed to justify BDM implementation may actually be not that weak.

assignment mechanism. We use the following notation to characterize a lottery $\mathcal{L}$

$$\mathcal{L} = \langle [(o_{11}, o_{21}, \ldots, o_{n1}), p_1]; [(o_{12}, o_{22}, \ldots, o_{n2}), p_2]; \ldots; [(o_{1m}, o_{2m}, \ldots, o_{nm}), p_m] \rangle$$

$\mathcal{L}$ is a set (inside the inner product sign) of prospects (each prospect is in square brackets) where o_{ij} is the i^{th} coordinate of the outcome vector (in round brackets) that would result if the j^{th} state realizes, which happens with subjective probability p_j; in case p is a density we will denote

$$\mathcal{L} = \left\langle \{[(o_j), \mathrm{d}p_j]\}_{j \in supp(p)} \right\rangle$$

the lottery which draws j according to p and assigns outcome (vector) o_j. Outcomes may themselves be lotteries (or set of lotteries), thus generating a compound lottery as standard in the literature.

Treatment is (potentially) a complicated object: after being treated agents may have to take some actions whose return is uncertain. What is important to realize (especially for when we will discuss potential failures of this seemingly robust class of mechanisms) is that the object auctioned-off is a set of effort-dependent subjective lotteries. Be as it may, the planner is a seller of this object.

Definition 2.8. A **treatment lottery** is a lottery L for which $o^L_{1,j} \in \{T, NT\}$, $o^L_{2,j} \in \mathbb{R}^-$ for each prospect j. This means that if j is drawn,[51] the holder of lottery L is assigned treatment status $o^L_{1,j}$ and has to pay a monetary sum $o^L_{2,j}$. The **treatment intensity** of a treatment lottery L is defined as

$$\tau(L) = \int_{\{j: o_{1,j} = T\}} \mathrm{d}p^L_j$$

Actions as selection intensities: the BDM class

Three assumptions on individual preferences are imposed.

Assumption MM Fix an outcome profile o and a positive number m. Then $\langle [o; 1] \rangle = \gamma_i (\{\langle [o; 1] \rangle, \langle [(o, -m); 1] \rangle\})$ where γ_i is the choice correspondence of in-

[51] Suppose for a moment that a fair device runs the treatment lottery, that is the lottery has j drawn with probability p^L_j announced by the planner and believed by all agents (this is going to be the validity constraint of the sampling procedure imposed in the next section).

dividual i and $(o,-m)$ be the outcome vector modified adding a new entry of m monetary expense.[52]

All agents prefer the treatment status, but will not pay an infinite amount of money to get treated.

Assumption PT For all i it holds $\langle [T;1] \rangle = \gamma_i \left(\langle [T;1] \rangle , \langle [NT;1] \rangle \right)$. There exists $\overline{u} \in \mathbb{R}$ such that $\langle [(NT,0)\,;1] \rangle = \gamma_i \left(\{ \langle [(T,-\overline{u})\,;1] \rangle , \langle [(NT,0)\,;1] \rangle \} \right)$ for all i.

Assumption AAM Suppose there are $o^{'}, o^{''}$ such that $\left\langle \left[o^{'};1 \right] \right\rangle \in \gamma_i \left(\left\langle \left[o^{'};1 \right] \right\rangle , \left\langle \left[o^{''};1 \right] \right\rangle \right)$. Consider a lottery $\mathcal{L}$ for which there is a prospect j with $o^{\mathcal{L}}_{\cdot,j} = o^{'}$, and let $\mathcal{L}^{'}$ be the modified lottery with $o^{\mathcal{L}'}_{\cdot,i} = o^{\mathcal{L}}_{\cdot,i}$ for $i \neq j$ and $o^{\mathcal{L}'}_{\cdot,j} = o''$. Then $\mathcal{L} \in \gamma_i \left(\mathcal{L}, \mathcal{L}' \right)$. If $\left\langle \left[o^{'};1 \right] \right\rangle = \gamma_i \left(\left\langle \left[o^{'};1 \right] \right\rangle , \left\langle \left[o^{''};1 \right] \right\rangle \right)$ and $p_j > 0$,[53] then $\mathcal{L} = \gamma_i \left(\mathcal{L}, \mathcal{L}' \right)$.

MM is a standard money monotonically assumption, agents prefer more money than less. PT states that all agents prefer the treatment status over the no treatment status.[54] AAM is a standard Anscombe and Aumann monotonicity saying that when we change some prospects of a lottery by putting better outcome we get an improvement in preferences. This assumption becomes critical once we allow outcomes to be other lotteries and, a fortiori, when as in treatment lotteries outcomes are pairs of monetary outflows and set of subjective lotteries. The discussion after the implementation result clarifies why this assumption is critical and how its failure will lead to loss in possible stimuli. AAM is however not necessary to have the following object well definite

Definition 2.9. For each individual i the **reservation price** $u^i \in [0, \overline{u}] = \overline{U}$ is given by

$$u^i = \sup_{u \geq 0} \left\{ u : \langle [(T,-u)\,;1] \rangle \in \gamma_i \left(\langle [(T,-u)\,;1] \rangle , \langle [(NT,0)\,;1] \rangle \right) \right\} \tag{2.16}$$

Definition 2.9 should be interpreted strictly: the reservation price is the maximum amount of money that individual i is willing (or able) to pay in order to change its environment from "non treated for sure" to "treated for sure".[55]

[52]Or subtracting it to existing monetary entries in the outcome vector.

[53]In case we have a lottery with a density, the condition should be replaced with if there is a set J of positive measure containing strictly better outcomes, then we have strict domination.

[54]Clearly it would be without loss of generality to say that there exists $\underline{u}$ such that $\langle [(T,\underline{u})\,;1] \rangle = \gamma_i \left(\{ \langle [(T,\underline{u})\,;1] \rangle , \langle [(NT,0)\,;1] \rangle \} \right)$. $\underline{u} = 0$ is just convenient.

[55] A theory gives statistical content to this empirical object by specifying the determinants of this reservation price and how they correlate with other determinants of success of the treatment. This, absent a utilitarian -value of treatment- interpretation of the reservation price may be particularly tricky (all examples in this paper use an utilitarian motivation).

The coercive stimuli we study are identified with set of treatment lotteries the planner quotes and from which each citizen can choose its most preferred.

Definition 2.10. A **Randomized Control Trial** mechanism with intensity x (RCT-x) has the government offer the singleton set of treatment lotteries

$$\Lambda^{RCT}(x) = \{\langle [T; x], [NT; 1-x] \rangle\}$$

Agents responding to an RCT will have no option but to take the unique lottery inside the RCT: they can only choose to be treated with probability x. At the opposite end we have an assignment mechanism that was already informally introduced in the vaccination example.

Definition 2.11. A **Conditionally Deterministic (CD)** mechanism with reservation u (CD-u) has the government offer the doubleton set of degenerate treatment lotteries

$$\Lambda_{CD}(u) = \{\langle [(T, -u); 1] \rangle ; \langle [(NT, 0); 1] \rangle\}$$

The term conditionally deterministic contrast with the following class, in which even conditional on their reservation price agents end up in a stochastic treatment assignment. This class is known as the class of Becker-DeGroot-Marschack mechanisms.[56]

Definition 2.12. Let $g : \overline{U} \to [0,1]$ be a weakly increasing function. A **Becker-DeGroot-Marschack mechanism** with selection intensity g (BDM$-g$) is the set

[56] BDM mechanisms were initially introduced to elicit the certainty equivalent of a monetary lottery. The agent held a lottery and he was asked to quote a price at which to "sell" it. Then a price was drawn from a distribution with full support over some interval that contained the prize state, and if the price drawn exceeded the one quoted by the agent, the lottery was sold and the agent earned the price quoted, on the contrary the agent kept the lottery without receiving any further monetary income. Use of BDM as a treatment assignment mechanism is proposed in CPS.

of treatment lotteries $\Lambda_g^{BDM} = \left\{\mathcal{L}_{g,u}^{BDM}\right\}_{u\in U}$ where for each u the lottery $\mathcal{L}_{g,u}^{BDM}$ is defined as

$$\mathcal{L}_{g,u}^{BDM} = \left\langle [(T,0)\,;g(0)]\,, \{[(T,-u')\,;\mathrm{d}g(u')]\}_{u'\in(0,u]}\,, [(NT,0)\,;1-g(u)]\right\rangle$$

Prospects are constructed by pairing treatment with a monetary outflow of u' for any $u' \leq u$, and pairing no treatment with no monetary expense for all $u' > u$. A price u' is then drawn according to g (seen as a CDF) and the associated treatment assignment and transfer is imposed.

The BDM mechanism is a non-competitive auction structure, as agents do not bid one another, but they place a bid against a fair randomization machine. Dominance of truthful revelation, that is to choose lottery $\mathcal{L}_{g,u}^{BDM}$ in Λ_g^{BDM} follows (see the proof of Proposition 2.3), from an argument that is similar to a second price auction in which we replace the source of strategic uncertainty with uncertainty on the fair draw from the lottery. However, the discussion after Proposition 2.3 discusses a branch of the literature on BDM mechanisms that highlights why one still needs to be careful in using the domination argument that imply BDM implementation in this environment.

It is immediate to notice that the RCT and CD assignment mechanisms, defined respectively in Definition 2.10 and 2.11 are just special cases of BDM mechanism: an RCT with intensity x has $g(u) = x$ for all $u \in \overline{U}$, while a CD with reservation price q is a BDM with $g(u) = \mathbf{1}_{u\geq q}$. Both RCT and CD mechanisms belong to the class of linear BDMs, that are those characterized by a selection intensity

$$g_{i,q}(u) = [\![iu+q]\!]_0^1$$

where $[\![x]\!]_\alpha^\beta = \max\{\min\{x,\alpha\}\,,\beta\}$.[57] For $i = 0$ we have the RCT with intensity q, while setting $q_i(u) = -\frac{u}{i}$, as $i \to \infty$ the pair $(i, q_i(u))$ approximates the CD with intensity u. In what follows we will denote the linear BDM (∞, u) to be such limit. In this sense RCT and CD correspond to limit points ($i = 0$ and $i = \infty$) of linear BDMs, formalizing the idea that BDM move smoothly between purely coercive and purely voluntary assignment mechanisms.

[57]Basically the outer function $[\![x]\!]_\alpha^\beta$ only guarantees that $g(u)$ is always a probability (bounded in $[0,1]$); one can actually use any transformation function $\Phi : \mathbb{R} \to [0,1]$, increasing and apply it to $iu+q$.

BDM implementation in superpopulation and potential failures

Given a stimulus in the form of a treatment lottery set, planner is interested in individual response in the form of the lottery chosen by each individual.

Definition 2.13. Fix a lottery domain Λ and a characteristic $u : X \to U$ that is minimally sufficient for treatment. The **policy correspondence** $L^{*}(\Lambda) : U \rightrightarrows \Lambda$ is given by[58]

$$L \in L^{*}(\Lambda)(u) \iff L \in \gamma_{u}(\Lambda)$$

A selection intensity $g : \overline{U} \to [0,1]$ is **implementable** (fully implementable) **in an individual market scheme** if there exists a set of lotteries Λ_{g} such that, for each $u \in U$, there exists a lottery (for every lottery) $L \in L^{*}(\Lambda_{g})(u)$ such that

$$\tau(L) = g(u)$$

Where $\tau(L)$ is the treatment intensity as in Definition 2.8. Notice that in principle the two sets $\overline{U}, U$ are different; the former is the set of reservation prices, while the latter is the sample space of a statistic that is minimally sufficient for treatment. The following proposition gives a version of the BDM result applied to this framework.

Proposition 2.3. *Under assumptions MM, PT, AAM, a selection intensity g is (fully) implementable if and only if g is (strictly) increasing, $g \in I(U, [0,1])$. In particular Λ_{g}^{BDM} implements g.*

From now on we will identify a stimulus with a weakly increasing selection intensity (and $\overline{U}$ with U) and keep implicit the treatment lottery construction. The following Remark is a version of Proposition 7 in CPS, which formalizes that $\overline{U} = U$ for price-incentivized assignment mechanisms.

Remark 2.1. Under assumptions MM, PT, AAM, if the set of stimuli coincides with BDM mechanisms, then the reservation price u of Definition 2.9 is minimally sufficient for treatment.

[58] By γ_{u} it is meant the choice correspondence of any individual with minimal sufficient statistic equal to u. The object is well definite by (minimal) sufficiency for treatment; under another characteristic s, γ_{s} would simply not be definite.

It is worth exploring potential failures of BDM implementation, as the objective is to characterize robustly individuals' response to a wide class of different selection policies by means of a low dimensional minimal sufficient characteristic. The approach seems promising as assumptions MM, PT, AAM are rather weak. However AAM becomes critical once we allow outcomes to be lotteries themselves, or even set of subjective lotteries as it implies a sort of indifference in the reduction of composite lotteries. The experimental literature (Grether and Plott (1979), Lichtenstein and Slovic (1971)) testing the functioning of the BDM mechanism documented the so called "preference reversal" phenomenon of having a monetary lottery X being preferred to a monetary lottery Y while price required to sell X was lower than the price required to sell Y. Karni and Safra (1987) show that, interpreting the BDM as a two stage monetary lottery, the preference reversal phenomenon can be explained by a failure of the independence axiom in the lottery reduction. For our purpose it is not important that the BDM does not elicit the true certainty equivalent of the treatment lottery, indeed the reservation price can be taken as an "empirical object" that determines response to the different treatment lotteries.[59] What will instead make Proposition 2.3 and Remark 2.1 fail is the fact that if such reduction does not hold, "results are not independent of the range of the announced prices": responses change as we change the interval from which the respondent has to choose. This would imply that the same individual, faced with BDM lotteries with different intensities could behave as if he had different reservation prices as defined in 2.9, so that each agent would be characterized by an higher dimensional type characterizing behavior conditional on all the possible policies inside the BDM intensities.[60] The fact that empirical work suggested a failure of BDM implementation even in the simple case where the lottery faced by agents was objective and monetary is somehow discouraging on the robustness of assumptions that lead to Proposition 2.3.

2.4.4 Super-population Propensity score (SP-PS) and Selection Regimes

The following object is popular in the literature of estimation of treatment effect

[59]If that was the case, the only further complication would come in the speculation model when conjecturing about the determinants of this characteristic.

[60]It would however by definition remain sufficient for all selection intensities in the conditionally deterministic regime.

Definition 2.14. Given $\omega \in \Omega$ and $g \in I(U, [0,1])$ the **super-population propensity score** (SP-PS) is given by

$$\tau(g,\omega) = \mathbb{E}_{\omega}[g(\mathbf{u})] = \int_U g(u)\,\mathrm{d}\omega_U(u)$$

Let $\{U_g(x)\}_{x\in[0,1]}$ be the upper contour sets of the function g, that is

$$U_g(x) = \{u \in U : g(u) \geq x\}$$

notice that $g \geq g'$ implies $U_g(x) \subseteq U_{g'}(x)$ for all x. Also, let $g^{-1} : [0,1] \to U$ be such that

$$U_g(x) = \left[g^{-1}(x), \overline{u}\right] \tag{2.17}$$

with the convention that $g^{-1}(x) = \overline{u}$ if $x > g(\overline{u})$.

I compute the SP-PS for familiar selection intensities

Example 2.3. SP-PS of RCTs is independent on the reduced form model, indeed

$$\tau(g_{0,x},\omega) = \int_U x\mathrm{d}\omega_U(u) = x$$

CD mechanism instead have

$$\tau(g_{\infty,u},\omega) = \int_u^{\overline{u}} \mathrm{d}\omega_U(u) = \omega_U([u,\overline{u}])$$

.

The following lemma will be useful in the next section.

Lemma 2.1. *For all action and social state pairs $\{g,\omega\} \in \mathcal{A} \times \Omega$, it holds*

$$\tau(g,\omega) = \int_0^1 \omega_U(U_g(x))\,\mathrm{d}x \tag{2.18}$$

where on the RHS we have a Riemann integral.

I then introduce classes of selection intensities (regimes) that are ranked according to an intensity index (higher intensity selections have each type participate to treatment with higher probability), and rich enough to have for each nondegenerate model a map from SP-PS to elements of the class which induce them. Formally,

Definition 2.15. A **regime** is a family of selection intensities $\mathcal{G} \subset I(U, [0,1])$ such that

i) $\mathcal{G} = (g_q)_{q \in \mathcal{Q}_\mathcal{G}}$ where the index set $\mathcal{Q}_\mathcal{G}$ is a compact subset of $\mathbb{R}$ such that

$$q > q' \implies g_q > g_{q'}$$

where $g > g'$ means $g(u) \geq g'(u)$ for all u and there exists a set of positive (Lebesgue) measure where inequality is strict.

ii) For all models ω such that ω_U has full support U, the mapping $\overline{\tau}^{-1}_{\mathcal{G},\omega} : [0,1] \to \mathcal{G}$ given implicitly by

$$\tau\left(\overline{\tau}^{-1}_{\mathcal{G},\omega}(x), \omega\right) = x \tag{2.19}$$

is well definite.

That is, for each model and SP-PS we can find a selection intensity inside the regime (hence, an index $q \in \mathcal{Q}_\mathcal{G}$) that obtains the desired SP-PS (joint with the model). We are already familiar with some regimes

Example 2.4. The set of RCTs constitute a regime under $\mathcal{Q}_\mathcal{G} = [0,1]$ and $\overline{\tau}^{-1}_{RCT,\omega}(x) = x$, independent of ω. CD intensities also constitute a regime under index set $\mathcal{Q}_\mathcal{G} = U$ (endowed with the reverse euclidean order to satisfy *i*)) and the function $\overline{\tau}^{-1}_{CD,\omega}$ is defined implicitly by

$$\omega_U\left(\left[\overline{\tau}^{-1}_{CD,\omega}(x), \overline{u}\right]\right) = x$$

The class of linear BDMs with fixed slope and nondegenerate transformation Φ also constitute a regime with $\overline{\tau}^{-1}_{i,\omega}$ defined implicitly by

$$\int_U \Phi\left(iu + \overline{\tau}^{-1}_{i,\omega}(x)\right) \mathrm{d}\omega_U(u) = x$$

Regimes are important as there is a sense in which selection inside a regime has a similar shape (from total coercion for an RCTs to total discretion in a CD), what changes is only the intensity represented by the index, an higher intensity making everyone more likely to choose treatment. Regimes as classes of assignment mechanisms will be used in the following discussion and in the discussion of the outcome models.

2.4.5 The Informational Content of an Assignment Mechanism

The key insight in CPS is that assignment mechanisms are source of information. As the final dataset contains stacked vector of individual messages, observing treatment relevant covariates will not just provide information on the marginal ω_U but on the whole correlation with potential outcomes and messages.

Definition 2.16. The **identification domain** of g, denoted $U_g \subset \mathcal{B}(U)$ is the coarsest $\sigma-$algebra under which g is measurable.

It is immediate to notice that $U_g = \sigma\left(\{U_g(x)\}_{x\in[0,1]}\right)$, and that $U_g = \mathcal{B}(U)$ if and only if g is strictly increasing in its domain.

Example 2.5. We have $U_{(0,q)} = \{\emptyset, U\}$, RCTs are not informative. For CD it holds[61] $U_{(\infty,u)} = \sigma(\{\emptyset, U, [0,u], [u,\overline{u}]\})$. In general a linear nondegenerate BDM has

$$U_{(i,q)} = \sigma\left(\left\{\left[0, \left[\!\left[\frac{-q}{i}\right]\!\right]_0^{\overline{u}}\right], \mathcal{B}\left(\left[\left[\!\left[\frac{-q}{i}\right]\!\right]_0^{\overline{u}}, \left[\!\left[\frac{1-q}{i}\right]\!\right]_0^{\overline{u}}\right]\right), \left[\left[\!\left[\frac{-q}{i}\right]\!\right]_0^{\overline{u}}, \overline{u}\right]\right\}\right)$$

The identification domain is an important object as it gives the information which is obtained in a selection mechanism that requires agents to report a message and assigns the treatment probabilities based on such messages, as the following Lemma formalizes.

Lemma 2.2. *If the selection intensity is implemented through a mechanism and the set of messages sent by individuals enters the information set of the planner jointly with outcome messages s, y, then after running the assignment mechanism with a*

[61] For $A, B, \ldots$ a sequence of sets, let $\sigma(A, B, \ldots)$ the $\sigma-$algebra generated by that sequence.

continuum of agents the distribution of random vector $\mathbb{E}^{\omega^{RF}}\left[\mathbf{u},\mathbf{y},\mathbf{s}\middle|U_g\right]$ *is estimated with precision.*[62]

The argument is as follows: by the revelation principle it is without loss of generality to focus on direct mechanisms, where agents report their reservation price and are assigned the lottery $\mathcal{L}_{g,u}$. With a continuum of agents the empirical distribution of reported types will estimate the reduced form with infinite precision. If $U_g = \mathcal{B}(U)$ then we observe the distribution of $\left[\mathbf{u},\mathbb{E}^{\omega^{RF}}\left[\mathbf{y},\mathbf{s}\middle|\mathbf{u}\right]\right]$ which contains the outcome model.

For any assignment function $g \in I(U,[0,1])$ and $\epsilon > 0$ there exists a strictly increasing function $\widetilde{g}$ such that $\|g-\widetilde{g}\|_\infty < \epsilon$. So we can estimate the whole outcome model[63] in one period by remaining arbitrarily close in norm to a "desired" selection intensity and therefore, by continuity of current payoff, suffering an arbitrarily small utility loss. The Markov kernels are not continuous in the selection intensity under the observational assumptions that messages sent in the assignment mechanisms are observed ex-post. Take as an example an interior RCT-x and a linear BDM with slope ϵ and intensity x: by Example 2.5 we have $U_{(\epsilon,x)} = \mathcal{B}(U)$ for $\epsilon > 0$ and small, while $U_{(0,x)} = \{\emptyset, U\}$. Informational discontinuities of this kind make the exploration-exploitation trade off disappear and reduce the dynamic problem to a sequence of static ones, with the exception of the first period in which optimal static policy has to be perturbed to obtain outcome model revelation.[64] In a simpler form, this insight was present even in the vaccination example: since all positive prices identify the model, as $N = \infty$ period decisions were statically optimal whenever $q^*(\mu) > 0$; if $q^*(\mu) = 0$ no solution existed as the "optimal" price is to set $q = \epsilon > 0$ as small as possible.[65] Finiteness of the population solve the estimation discontinuities as those in the vaccination example but not the identification discontinuities: the BDM-ϵ is "discontinuously more informative" than an RCT, as the former produces a dataset of the type $\{u_i\}_{i=1}^N$ and the latter $\emptyset$. We don't like information discontinuity in selection mechanisms, as they make the information acquisition problem trivial.

[62]This is not completely accurate as we would need to guarantee that $g > 0$ otherwise on the atom of the partition associated to $g = 0$ no one wold be treated and we could not observe potential outcomes and messages.

[63]We cannot estimate the whole reduced form model as some correlation between s and y is missing. However recall that for static optimality only the outcome model is necessary.

[64]The solution to the dynamic problem would prescribe to play in the first period a small perturbation of the static optimal mechanism, and from the second period onward the objectively optimal mechanism under the revealed state.

[65]In general, information discontinuities create similar existence failures.

In next paragraph discuss possible remedies. The sampling procedure we discuss in Section 2.5 can implement both implicit and explicit BDMs; Proposition 2.5-i) will prove that under observation of the empirical propensity score, implicit mechanism preserve continuity of the Markov kernels in the selection intensities. For those reasons, implicit implementation is discussed first and will be the focus of the remainder of this chapter.

Implicit BDMs

The fact that an increasing function is implementable by a (direct) mechanism does not imply that such mechanism is actually used or that messages sent inside this mechanisms are saved by the sampling procedure and transmitted to the planner who can process it as information about the reservation prices.

Definition 2.17. A treatment choice set C_u is a pair of treatment lotteries

$$C_u = \{\langle [(T, -u)\,; 1] \rangle\,, \langle [(NT, 0)\,; 1] \rangle\}$$

An **implicit BDM mechanism with intensity** g is a single lottery with outcomes treatment choice sets, given by

$$\mathcal{IL}_g = \left\langle [C_0; g\,(0)]\,, [C_u; \mathrm{d}g\,(u)]_{u \in U^o}\,, [C_{\overline{u}}; 1 - g\,(\overline{u})] \right\rangle$$

Example 2.6. In implicit form, a CD mechanisms is (after removing zero probability prospects)

$$\mathcal{IL}_{CD_u} = \langle [C_u; 1] \rangle$$

while an RCT is

$$\mathcal{IL}_{RCT_x} = \langle [C_0; x]\,, [C_{\overline{u}+\epsilon}; 1 - x] \rangle$$

Without any signal being sent the implicit mechanism assigns each citizen a choice set drawn from $\mathcal{IL}_g$. The citizen then chooses between the treatment lotteries in the drawn choice set and the planner gets to observe only the realized fraction of

IMPLICIT BDM-g	EXPLICIT BDM-g
Do not send any message	Choose $\mathcal{L}_{g,u} \in \Lambda_g^{BDM}$ by revealing the type
Choice set C_u is drawn from g	Mechanism draws a prospect from $\mathcal{L}_{g,u}$
Citizen chooses $\langle (T, -u) ; 1 \rangle$ or $\langle (NT, 0) ; 1 \rangle$	Citizen gets the outcomes from the prospect
Treatment status is revealed	Treatment status is revealed (but irrelevant)

Table 2.1: Implicit vs Explicit mechanisms (in super-population)

agents that end up treated. Difference between implicit and explicit mechanisms are presented in Table 2.1. Under mild assumptions[66] it holds $\mathcal{IL}_g \sim^u \mathcal{L}_{g,u}$, indeed the latter is the lottery that obtains substituting the lottery that gets chosen by type u at each choice set offered in the implicit mechanism.[67] Implicit and explicit mechanism induce therefore the same stochastic map from population to sample (see Section 2.5, the sampling procedure takes only sampling intensity and not the implicit-explicit type of mechanism). The difference is however important once we consider that the reported type through which explicit mechanism assign the lottery can be used as messages to make inference: if the assignment mechanism is implicit only the mean $\tau(g, \omega)$ is identified which, contrary to the explicit homolog, does not characterize ω_U.[68]

Alternative failures of information discontinuity

Quasilinear mechanisms with imperfect signal Agents choose a lottery that place them into treatment with probability $g(u)$ and pay (upfront) transfer $t(u) = \int_0^u u' \mathrm{d}gu'$. The market records the sale for a $g(u)$ lottery with an error ϵ so that the planner receives a message $m = g(u) + \epsilon$ for $\epsilon \sim (0, \nu)$ with ν known positive but small. By reducing the slope of the BDM the planner is effectively increasing the importance of the noise (ϵ is distributed independently of the g function). Suppose that the distribution of reservation prices is $\mathcal{N}(\mu, 1)$, ϵ also normal and focus on linear BDMs with $g(u) = q + au$. With infinite samples the slope of the mechanism

[66] Basically we require there to be an indifference on flexibility, but this is immaterial as there is no real time passing from choice of a lottery and realization of its outcome.

[67] As $\langle [(T, -u) ; 1] \rangle \in \gamma_u(C_{u'}) \iff u' \leq u$, then $\mathcal{IL}_g^*(u) = \mathcal{L}_{g,u}$.

[68] An implicit mechanism gives a a garbled signal of the associated explicit one, indeed the true treatment status conditional on the reservation price is just obtained through a further randomization.

is irrelevant (provided it is positive) and accurate estimation of the nonparametric distribution is always possible.[69] As

$$Lik(s_i, \mu) \propto \frac{1}{(a^2 + \nu)^N} \exp\left\{\sum_{i=1}^{N} (s_i - (q + a\mu))^2\right\}$$

the statistical (Fisher) information[70] contained in a single sample point is

$$I(\mu) = E\left[\frac{\partial^2}{\partial\mu^2} ln\left[Lik(s, \mu)\right]\right] \propto -\frac{2a}{(a^2 + \nu)}$$

Therefore for any fixed $\nu > 0$ as $a \to 0$ each observation becomes uninformative as the noise is overwhelming, approaching continuously the uninformative limit $a = 0$. This construction gives the most clean mathematical representation of information continuity, however the assumptions behind it are somehow ad-hoc: first of all the revelation principle is abandoned by having the message be a noisy observation of the selection probability $g(u)$ rather than of u; if agents were to report u and then assigned $g(u)$, changing the slope of the latter would have no effects on information quality. Secondly, there is no clear economic interpretation for the recording error ϵ.

Partitional BDMs Prices may be restricted to lie in a discrete grid, i.e. g has to be a step function.[71]

From a estimation perspective, it is interesting the case in which we can choose a partition $a_1 = 0 < a_2 < \cdots < a_N = \overline{u}$ that categorizes our observations (that is for each citizen we can tell ex post what atom of the partition he belongs to). In case we have ω_U belong to a parametric family, such partition will give a set of moment conditions

$$\{F_\theta[a_{i-1}, a_i]\}_{i=2}^{N}$$

that we can exploit for estimation.[72]

Partitional BDMs may be relevant for the present application as under the sampling procedure used to implement the intensity function in a finite population, agents

[69] As ϵ has known distribution, if we had an infinite sample, then we would be able to recover the distribution of u as difference of random variables with known distribution.

[70] Which is relevant as by the Cramer-Rao inequality its inverse bounds from below the variance of an unbiased estimation of μ

[71] There is a natural sense in which prices are discrete, and moreover we can imagine some costs of increasing the number of treatment lotteries quoted on the market.

[72] A natural way to proceed it to exploit the known formulas for the asymptotic variance of the GMM estimator to make an approximate (this approximation is actually non-trivial as we need to make all posterior beliefs “closed” to those that would obtain under likelihood updating) optimal choice of the partition.

end up to be treated with probabilities $g\left(\frac{k}{N}\right)$ $k \in [N]$: basically we partition the g function vertically with steps of size $\frac{1}{N}$.

Parametric models If we had parametric uncertainty, namely our reduced form model took the form $\mathrm{d}\omega_U(u) = \mathrm{d}\omega_U(\theta, u)$ for $\theta \in \Theta$ finite dimensional (for example, take the vaccination case in which γ, the parameter of the demand function with known functional form), then the identification requirement is much less demanding. Consider indeed a CD mechanism with reservation price u which under the nonparametric version had no identification power; if $F_\theta(u) \neq F_{\theta'}(u)\ \forall \theta, \theta' \in \Theta$, then θ is identified. We would however have discontinuity moving away from intensities in the RCT regime.

We do not expand further on those arguments, but we do stress again that the selection properties of a sampling intensity are not necessarily related to its estimation part, and in particular that guaranteeing truthful revelation does not necessarily imply that planner gets to known the reservation price. We can now study how the selection intensity is implemented in the finite population and offer an implementing procedure that is flexible enough to have both implicit and explicit mechanisms.

2.5 Population Implementation of the selection scheme

So far, we only discussed the superpopulation properties of selection mechanisms. At each period of time, a random draw of N individuals from the superpopulation distribution consititutes the population, the set of citizen responding to the current period stimulus and determining outcome (next section) and information collected.

A population $P \in \mathcal{P}$ is an ordered collection of N individuals, each characterized by a vector of reduced form characteristics (u, y, s). An individual inside a population will be called citizen. The population realizes as a sequence of N independent random draws from the super-population, so that a reduced form model induces a distribution over possible realized populations as sequences of types that is $\omega^N \in \Delta\left(\left(\Omega^{RF}\right)^N\right)$ is the product measure $\omega \otimes \cdots \otimes \omega$. Inside a population, citizens are identified with a label $i \in [N]$, $(u, y, s)_{i_P}$ denoting the reduced form type associated to citizen i in population P. After being drawn individuals must be divided in treatment and no treatment groups.

Definition 2.18. An $N-$**sampling procedure** is a map $\Sigma_N : I(U, [0,1]) \times \mathcal{P} \to \Delta\left(\{0,1\}^N\right)$; it takes a pair of population and assignment intensity and returns a *joint* probability distribution over sampling outcomes, 1 denoting treatment, 0 no treatment.

The procedure is **$\epsilon-$valid** if for all $u \in U$ $g \in I(U,[0,1])$, $i \in [N]$ and $P \in \mathcal{P}$ with $u_{i_P} = u$, it satisfies[73]

$$|marg_i[\Sigma_N(g,P)] - g(u)| \leq \epsilon \tag{2.20}$$

$0-$validity will be referred to as validity.

The ($\epsilon-$) validity condition is a minimal consistency assumption that requires the sampling procedure to respect what the selection intensity prescribes. It is immediate to notice that a valid sampling procedure is permutation invariant. The simplest valid procedure is the so called Bernoulli trial.

Example 2.7. (Bernoulli trial) Independently across citizens, the sampling procedure draws treatment status from a $Ber(g(u))$ random variable. Formally, for any P, g and any cylinder $A = \prod A_i$, $A_i \in \mathcal{P}(\{0,1\})$ it holds

$$\Sigma_N(g,P)(A) = \prod_{i=1}^{N} g(u_{i_P})^{\mathbf{1}_{1\in A_i}} (1 - g(u_{i_P}))^{\mathbf{1}_{0\in A_i}}$$

For this example not to mislead, notice that in general it is *not* required that Σ_N is a product measure: selection result can be correlated across agents, provided it is done in a type-independent manner. After the procedure is implemented, a sample realizes. The name is suggestive of the fact that inference about the outcome type is possible only using the set of individual that get treated.

Definition 2.19. The **sample** ς is the subset of the population which ends up receiving treatment status 1. The **empirical propensity score** $\bar{\varsigma} = \frac{1}{N}|\varsigma|$ is the ratio of the (stochastic) cardinality of the sample and the cardinality of the population (fixed at N).

[73]Notice that $marg_i[S_N(g,P)]$ is a measure over $\{0,1\}$, therefore with an abuse of notation I let $marg_i[S_N(g,P)] = (marg_i[S_N(g,P)])(1)$, the probability it assigns to being treated.

Conditional on the population P the sample is a random variable with realization a random subset of P and whose distribution depends on the selection intensity g and on the sampling procedure Σ_N. For each $v \in \{0,1\}^N$ let vP be the subset of P taking only citizens associated to indices which have value 1 in v, that is $vP = \left\{(u, y, s)_{i_P}\right\}_{\{i:v_i=1\}}$. Conditional on P, ς can only take the form $\varsigma = vP$, and

$$\mathbb{P}_{\Sigma_N}\left(\varsigma \middle| P, g\right) = \Sigma_N\left(g, P\right)\left(v\right) \text{ for } \varsigma = vP$$

Uncertainty on the realized sample integrates the population conditional likelihood derived above over likelihood of populations implied by the reduced form model. It is therefore a random variable with realization random subsets of $N-$dimensional random draws from the reduced form model.

$$\mathbb{P}_{\Sigma_N}\left(\varsigma \middle| \omega, g\right) = \int_{\mathcal{P}} \mathbb{P}_{\Sigma_N}\left(\varsigma \middle| P, g\right) \mathrm{d}\omega^N\left(P\right) \tag{2.21}$$

Notice that $\mathbb{P}_{\Sigma_N,\omega,g}\left(\varsigma\right) = \mathbb{P}_{\Sigma_N,\omega,g}\left(\varsigma'\right)$ for every ς' permutation of ς (this is due to permutation invariance of Σ_N and product measure for population). Given the sample, the properties of the empirical propensity score are obtained by applying the cardinality operator on the random variable.

Now let's reconsider the Bernoulli trial defined in Example 2.7. That procedure has clearly the flavor of a direct mechanism: each citizen in the population is required to send a message $u \in U$, and the sampling procedure *independently* assigns treatment status (and eventual monetary transfers) using a fair randomization device to implement the chosen selection intensity function (that is, draw treatment and payments pairs according to $\mathcal{L}_{g,u}$). Despite being valid, this procedure has a major drawback, as it induces too much uncertainty on the empirical propensity score, hence on the realized size of the treatment group. To visualize this, take the example of an RCT x where $x = \frac{M}{N}$ for some integer number $M < N$. Suppose we use a Bernoulli trial; it is immediate to see that the resulting sample size is distributed according to $Bin\left(x, N\right)$. This result holds more in general

Remark 2.2. Under a Bernoulli trial, $\overline{\varsigma}\left(g, \omega\right) \sim Bin\left(\tau\left(g, \omega\right), N\right)$.

However, the planner can easily conceive of an alternative valid procedure to implement an RCT which results in zero variance (of $\overline{\varsigma}$): pick a random subset of cardinality M of citizens and allocate them to treatment.[74] Notice the procedure

[74]Completely randomized experiment, in the language of Imbens and Rubin (2015).

remains valid: each citizen gets indeed treated with probability x as the likelihood he is drawn in the random subset of M elements is exactly $\frac{M}{N}$. What we break is independence across each one's actual treatment, which is *not* a requirement in our definition of validity. Kasy (2013) makes the point that, to minimize the variance of the treatment effect estimator, experimenters should not randomize over the sample size conditional on the covariates: in this framework this means that, rather than drawing the sample size from a distribution with a given mean by running Bernoulli trials one should choose a sample size then randomize on the citizen inside the fixed size (in an independent way so to maintain validity). Implementation of this principle is immediate for RCT regimes. However it is less obvious how one can reduce sample variability of the empirical propensity score in a generic stochastic assignment mechanism as the BDM, while preserving validity.

We propose the following sampling procedure, a graphical description of which is offered in Figure . From now on, Σ_N will denote this sampling procedure, expectations and probabilities being defined for this example.

Sampling procedure

- At the same time he is drawn in the population, the citizen is also assigned a label k, that defines the point $\frac{k}{N}$ on the domain $[0, 1]$ he is assigned to. Each label $k \in [N]$ is assigned to one and only one citizen: the mechanism guarantees that the *empirical* distribution of labels is uniform.[75]

- Citizen is offered treatment at price $g^{-1}\left(\frac{k}{N}\right)$, where $g^{-1} : [0, 1] \to U$ is defined in (2.17). In the figure, g^{-1} is the black function separating red and green regions.

- He compares this offered price with his reservation price (represented on the vertical line): if it is below, treatment is accepted and the price $g^{-1}\left(\frac{k}{N}\right)$ paid (red horizontal dotted line); else, he selects no treatment. Treatment status for each citizen is a function of both its reservation price and the random label he is assigned to

$$\widehat{\tau}_g(u, k) = \begin{cases} 1 & \text{if } g(u) \geq \frac{k}{N} \\ 0 & \text{else} \end{cases} \tag{2.22}$$

- If treated, citizen ends up with health status he carried from the super-population, and sends the signals (a set that he also carried from the super-population) that planner action induce.

[75] A simple way to do it is to assign each citizen their order in the drawing phase.

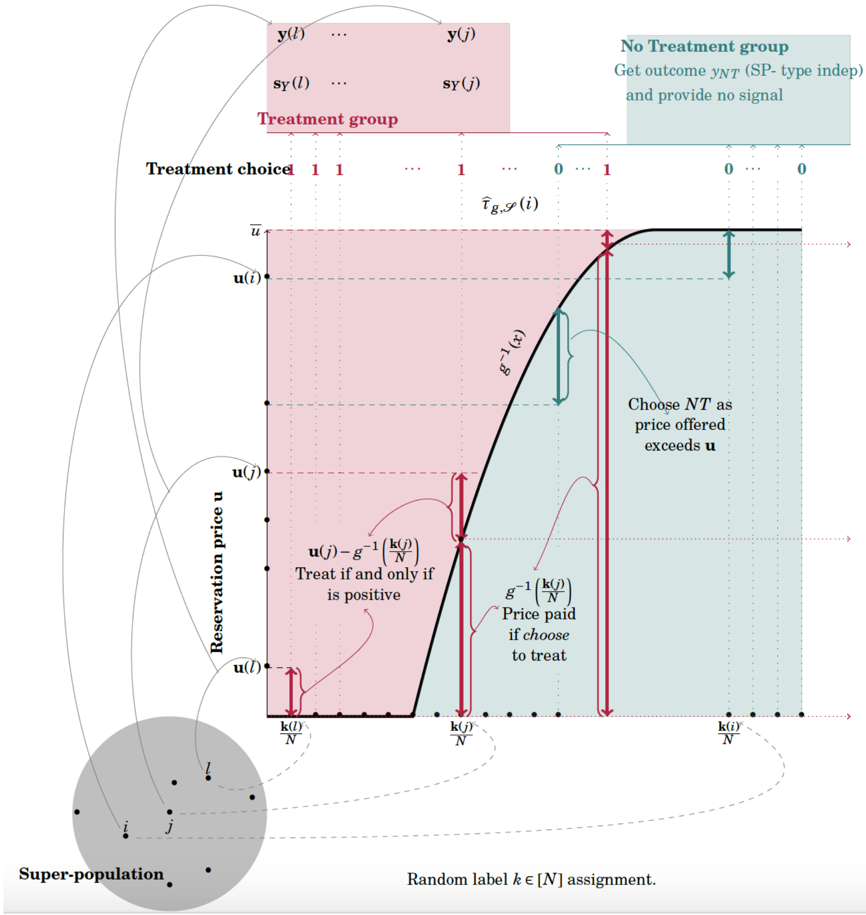

Figure 2.3: Population implementation of an implicit BDM

The sampling procedure implements the implicit BDM (see Table 2.1) in a correlated way across citizens. The Bernoulli trial associated to independent implementation of the BDM would have each citizen assigned randomly a label, disregarding the constraint on the empirical distribution. The explicit BDM would have the citizen report immediately his reservation price, then be assigned a label and have the mechanism make message-consistent choice between paying for treatment and staying untreated. Similarly, it can be implemented in a correlated or independent way across citizens. The following Proposition establishes some immediate properties of the proposed sampling procedure.

Proposition 2.4. *i) If* $g(u) \in \{0,1\}$, $\mathbb{P}_{\Sigma_N}\{\widehat{\tau}_g(u,\mathbf{k}) = g(u)\} = 1$*; the sampling procedure has no impact on conditionally deterministic selection intensities*

ii) The sampling procedure is $\frac{1}{N}-$*valid.*

Using this sampling procedure induced a (bounded) reduction in validity as compared to the Bernoulli trial. Basically we see that the likelihoods have discrete jumps of size $\frac{1}{N}$ (and so do prices in the partitional BDM that results). Why should we accept $\frac{1}{N}-$validity if we have a simple $0-$valid assignment mechanism as the Bernoulli trial? The following proposition clarifies the nice properties of this implementation mechanism in terms of the distribution it induces over the empirical propensity score $\overline{\varsigma}(g,\omega)$.

Proposition 2.5. *It holds*

i) The distribution of $\overline{\varsigma}(g,\omega)$ *is continuous in* g,ω

ii) $\left|\mathbb{E}_{\Sigma_N}\overline{\varsigma}(g,\omega)-\tau(g,\omega)\right|\leq\frac{1}{N}$ *for all* N,ω,g. *Hence* $\overline{\varsigma}(g,\omega)\xrightarrow{p}\tau(g,\omega)$.

iii) If $Im(g)\subset[a,b]\subset[0,1]$, *then* $\mathbb{P}_{\Sigma_N}\left(\left\{\overline{\varsigma}(g,\omega)\in\left[\lfloor a\rfloor_N,\lfloor b\rfloor_N\right]\right\}\right)=1$ *for any* $\omega\in\Omega$, *where* $\lfloor x\rfloor_N:=\max_{j\in\{0\}\cup[N]}:\frac{j}{N}\leq x$.

iv) For an RCT $g_{0,x}$, $\overline{\varsigma}(g_{0,x},\omega)\sim\delta_{\lfloor x\rfloor_N}$ *for all* ω.

v) Under a conditionally deterministic regime it holds $\overline{\varsigma}(g_{\infty,u},\omega)\sim Bin\left(\tau\left(g_{\infty,u},\omega\right),N\right)$.

Point $i)$ is the fundamental information continuity if only outcome was observable. Point $iii)$ is an important implication of the sampling procedure and it is going to be relevant in case there are "budget constraints", namely the amount of treatment units that can be provided each period is limited to a number $M<N$. In this case by having $g(\overline{u})<\frac{M}{N}$ one can guarantee that in no case the sample size (number of treatments provided) exceeds M. Intuitively, the procedure tries to remain as concentrated as possible around the targeted empirical propensity score (the SP-PS), but it has to satisfy individual validity constraint. Clearly, there is nothing to do under a conditionally deterministic regime to reduce sampling variability while keeping validity as each citizen will request a degenerate lottery.

The two figures below display how the sampling procedure determines the distribution of the empirical propensity score.

In Figure 2.4 we firstly have the planner choose an increasing selection intensity (top panel). This function is then inverted (second panel, with y scale on the right). The true reduced form model then, through its marginal ω_U induces a decreasing self-map on $[0,1]$ giving the values $\left\{\omega_U\left(\left[g^{-1}(x),\overline{u}\right]\right)\right\}_{x\in[0,1]}$, represented on the second panel with left y scale for three different models. The bottom panel takes the true model to determine the parameters of the parameters of the Poisson-Binomial EPS as $\left[\left\{\omega\left(U_g\left(\frac{k}{N}\right)\right)\right\}_{k\in[N]}\right]$.

Figure 2.5 instead plots in the top panel the distribution of empirical propensity scores for different selection intensities and models calculated so to induce the same SP-PS. Notice that, by Remark 2.2, a Bernoulli trial would have all those EPS be a Binomial $(SP - PS, N)$. The top panel illustrates points $i) - iii) - iv)$ of Proposition 2.5: the RCT induces a degenerate distribution for the EPS, while the CD conforms with the Binomial. Looking at the BDMs with different slope we see that as it gets steeper (approaching the RCT), the distribution becomes more concentrated around the SP-PS, while it is more dispersed when the empirical selection intensity is closer to a CD. This is exactly the crucial information continuity property $i)$. The bottom panel displays EPS distribution for a given selection intensity, but letting the models (and hence the SP-PS, I took truncated exponential distributions with different γ parameters) vary.

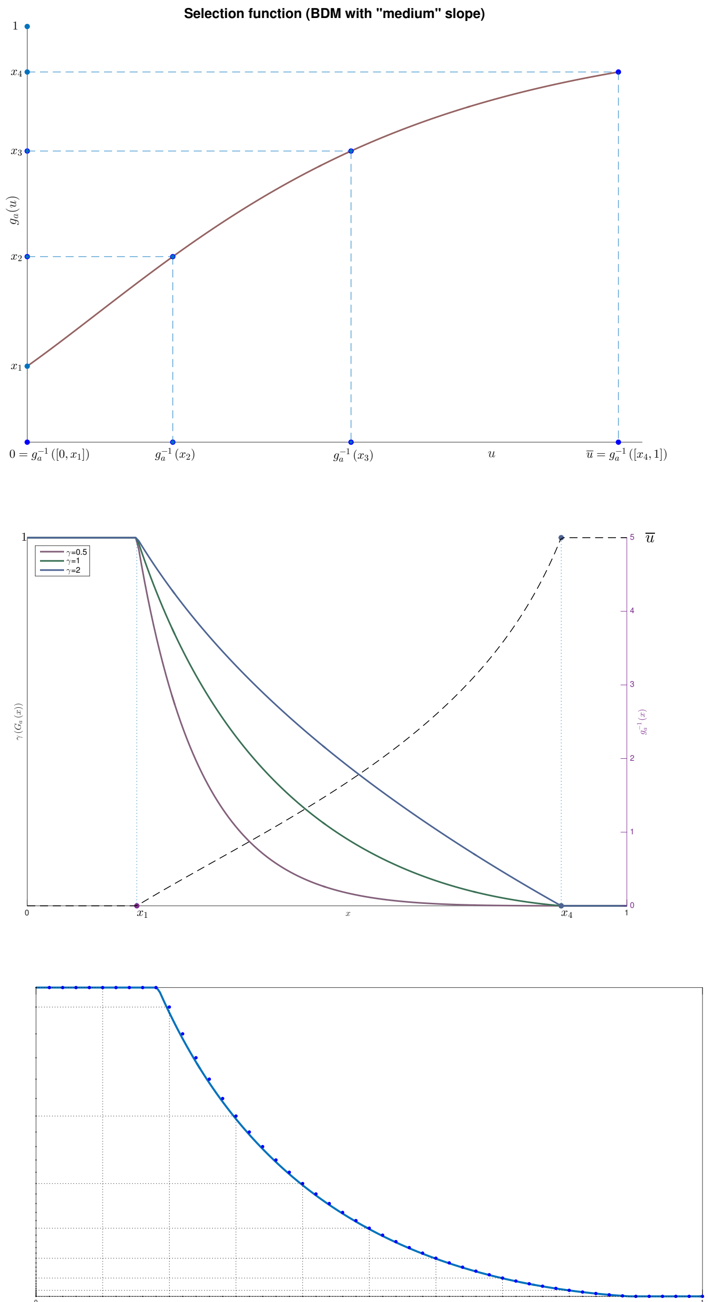

Figure 2.4: Top panel: selection intensity. Second panel: inverse selection intensity (right scale), $\omega_U(U_g(x))$ for different models (left scale). Bottom panel: parameters of the Poisson-Binomial.

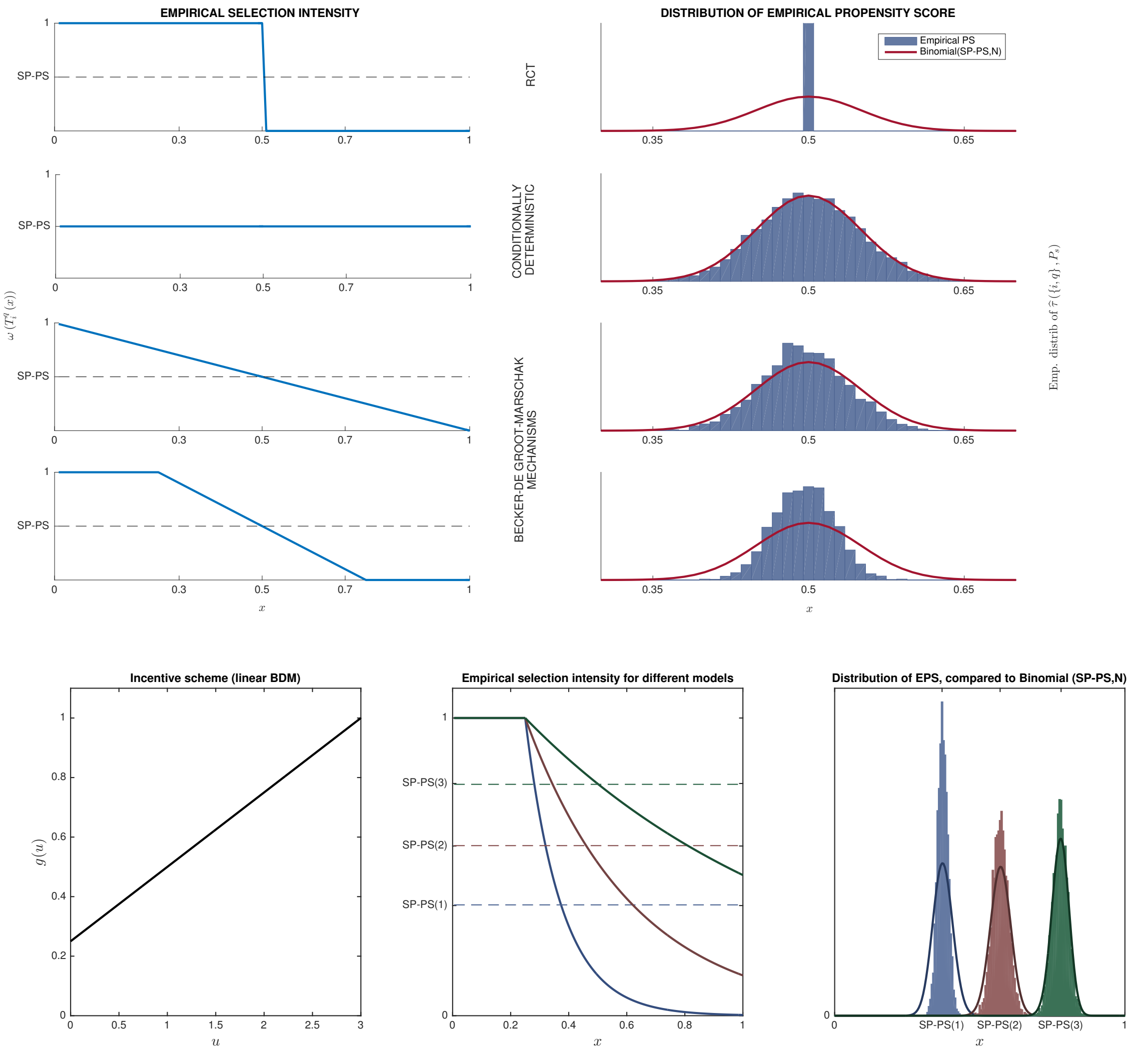

Figure 2.5: Distribution of EPS. Top panel: Different pairs of models and selection intensities having the same SP-PS. Bottom panel: Same selection intensities and different models.

2.6 Outcome model

The treatment assignment mechanism splits the population in the treatment group (sample) and a no treatment group. Citizens in the latter get a non-stochastic outcome, assumed to be 0. Citizens in the sample end up in an health status corresponding to their potential outcome. The speculation structure specifies determinants of this reduced form type (potential outcome, see Example 2.2) which drives the correlation between the incentivizable characteristic and the (ex-post) observables. This Section presents the outcome model, which defines the control properties of a selection mechanism and finally allows us to setup the problem of monetary incentivation into treatment as a MAB problem of the type presented in Section 2. Before doing that, we discuss how standard objects in the literature of estimation of treatment effect (average treatment effect, unconfoundedness ...) are defined in the setting presented in this Chapter. We define a novel object, the regime distortion function which tilts the RCT (identity) map from the propensity score into the average outcome conditional on treatment and use it to characterize unconfounded regime and model pairs. As for the control properties of selection mechanisms, we establish in Section 2.6.3 that CD mechanisms are control optimal under linear social preferences. In Proposition 2.5 we established that valid implementation of CD mechanisms induces the largest variance of the EPS among BDM mechanisms (an undesirable property per Kasy (2013)). Those two results provide therefore a novel experimentation-exploitation tradeoff channel which is peculiar to the application studied in this chapter.

2.6.1 Success Rate and Regime Distortion

Definition 2.20. Take a reduced form model ω as given.[76] The **average outcome conditional on treatment (ATO)** is given by

$$\sigma_\omega = \mathbb{E}_\omega [\mathbf{y}]$$

Given a reduced form characteristic s, the **conditional average treatment outcome (CATO)** is a function $\sigma_\omega : S \to [0, 1]$ given by

$$\sigma_\omega (s) = \mathbb{E}_\omega \left[\mathbf{y} \middle| \mathbf{s} = s\right]$$

We call s the trivial characteristic if S is a singleton.

[76] ATO, CATO are functions from Ω^{RF} to $[0, 1]$.

Standard in the literature is the name average treatment effect, we keep it different to stress the assumption that no treatment is a non-stochastic region. However the definitions coincide as $Y(0,i)=0$ for all i.[77] It is convenient to use the decomposition (2.14) to write

$$\sigma_\omega = \int_{U\times S} \sigma_\omega(u,s)\,\omega_{U\times S}\left(\mathrm{d}s\middle|\mathrm{d}u\right) = \int_S \sigma_\omega(s)\,\omega_S(\mathrm{d}s)$$

When evaluation the success rate of a given selection intensity however, one cannot simply look at ATO-CATO, because of selection effects.

Definition 2.21. Given $\omega\in\Omega$ and $g\in I(U,[0,1])$ the **super-population success rate** (SP-SR) is given by

$$\sigma(\omega,g) = \mathbb{E}_\omega\left[g(\mathbf{u})\cdot\mathbf{y}\right] = \int_U \sigma_\omega(u)\,g(u)\,\omega_U(\mathrm{d}u) \tag{2.23}$$

The superpopulation success rate the fundamental object for our analysis as it drives control optimality. It is only a function of the outcome model $\omega^O\in\Delta(Y\times U)$ but, as we will discuss shortly, equivalent characterizations in terms of unconfounded characteristics may simplify the conjecturing exercise taking into account the fact that U is a potentially large interval and we would need to estimate CATOs at all its points (see Proposition 2.8). Each reduced form model ω induces a map from stimuli to SP-PR that are exactly the correlated stochastic returns of policies in the multiarmed bandit interpretation of policymaking.

Selection effects are described by the distortion function, which measures the excess success rate of a particular selection intensity compared to the coercive regime achieving the same SP-PS. Inside a regime the distortion is a function of the intensities.

Definition 2.22. For a pair of model and selection intensities, the **distortion** $\phi:\Omega\times I(U,[0,1])\to\mathbb{R}$ is given by

$$\phi(g,\omega) = \frac{\sigma(\omega,g)}{\sigma_\omega\tau(\omega,g)} = \frac{\int_U \sigma_\omega(u)\,g(u)\,\mathrm{d}\omega_U(u)}{\sigma_\omega\int_U g(u)\,\mathrm{d}\omega_U(u)}$$

[77] Or, in case each no treated gets outcome 1 with probability σ_{NT}, then $ATE=\sigma_\omega-\sigma_{NT}$.

where τ is the super-population propensity score. Given a regime $\mathcal{G}$, the **regime distortion function** $\phi_{\mathcal{G}} : \Omega \to \mathbb{R}^{[0,1]}$ associates to each model a function[78]

$$\phi_{\mathcal{G}}(\omega) : [0,1] \longmapsto \mathbb{R}$$

$$\phi_{\mathcal{G}}(\omega)(x) = \phi\left(\omega, \tau_{\mathcal{G},\omega}^{-1}(x)\right) = \frac{\int_U \tau_{\mathcal{G},\omega}^{-1}(x)(u)\,\sigma_\omega(u)\,\mathrm{d}\omega_U(u)}{\sigma_\omega x}$$

where $\tau_{\mathcal{G},\omega}^{-1}(x) \in \mathcal{G}$, defined in (2.19) is the selection intensity belonging to regime $\mathcal{G}$ that under model ω would induce a SP-PS equal to x.[79]

Now let $\mathcal{G}_{RCT}$ be the regime of RCTs defined in Example 2.4 and notice that

$$\phi_{\mathcal{G}_{RCT}}(\omega)(x) = \frac{\int_U \tau_{\mathcal{G}_{RCT},\omega}^{-1}(x)(u)\,\sigma_\omega(u)\,\mathrm{d}\omega_U(u)}{\sigma_\omega x} = \frac{\int_U x\sigma_\omega(u)\,\mathrm{d}\omega_U(u)}{\sigma_\omega x} = 1$$

the regime distortion function is degenerate at 1. Also, notice that for *any* regime it must hold

$$\lim_{x\to 1} \phi_{\mathcal{G}}(\omega)(x) = 1$$

Now fix a regime $\mathcal{G}$ and a model ω and let for $q \in \mathcal{Q}_{\mathcal{G}}$, $\sigma_{\mathcal{G},\omega}(q) = \sigma(\omega, g_q)$. By definition it holds

$$\sigma_{\mathcal{G},\omega}(q) = \sigma_\omega \tau_{\mathcal{G},\omega}(q)\,\phi_{\mathcal{G},\omega}(\tau_{\mathcal{G},\omega}(q))$$

we can differentiate this function with respect to the intensity of the mechanism. Dropping subscripts we obtain

$$\sigma'(q) = \sigma\tau'(q)\left[\phi(\tau(q)) + \tau(q)\,\phi'(\tau(q))\right] \tag{2.24}$$

which gives the response in average outcome as we increase the intensity of the regime. Equation (2.24) gives an expression for the marginal change in the success rate due to increasing the intensity of selection[80] *while remaining inside the same*

[78]The regime distortion function answer the following question: what is the extra success rate under model ω if a fraction x enters treatment as a response to a stimulus belonging to the regime compared to success rate under the same model ω if instead the fraction x was randomly taken from the super-population and assigned treatment.

[79]Notice indeed $\tau\left(\omega, \tau_{\mathcal{G},\omega}^{-1}(x)\right) = x$ by definition.

[80]Recall that $q > q' \implies g_q > g_{q'}$ so everyone is more likely to be treated under an higher intensity selection.

regime. Same regimes are associated to similar estimation properties, therefore we can see incentive provision as the choice of a regime (for estimation purposes) and an associated intensity. The following example calculates a non-trivial distortion function.

Example 2.8. All agents know and it is true that if the get treated they survive with probability is e,[81] where e is the effort chosen $e \in [0,1]$. $k \in [0,1]$ is an effort disutility drawn from an unknown distribution characterizing the reduced form model. Conditional on being treated, agent of type k solves

$$\max_{e\in[0,1]} 2\left[e - \frac{1}{2k}e^2\right]$$

Which gives policy $e^*(k) = k$ and value $V(k) = k$ that is also the reservation price. The ATO is

$$\sigma_\omega = \int_0^1 e^*(k)\,\mathrm{d}\omega(k) = \int_0^1 k\mathrm{d}\omega(k)$$

Planner uses the conditionally deterministic regime: intensities $g_{\infty,q}$ are enforced by quoting price $q \in \mathcal{Q}_{\mathcal{G}} = [0,1]$ (recall for CDs the inverse order on the index set is used). Only types $k > q$ will select into treatment. Now we calculate the distortion of the CD regimes under the uniform model, that is assuming $k \sim \mathcal{U}[0,1]$ so that $\sigma_\omega = \frac{1}{2}$. To induce a fraction x to get into treatment the planner must charge a price[82] $\overline{\tau}^{-1}_{\mathcal{G},\omega}(x)$ that solves

$$Pr\left[k > \overline{\tau}^{-1}_{\mathcal{G},\omega}(x)\right] = 1 - \overline{\tau}^{-1}_{\mathcal{G},\omega}(x) = x$$

that is, $\overline{\tau}^{-1}_{\mathcal{G},\omega}(x) = 1 - x$. Now, the fraction of healthy citizens in the incentivized mechanism if x get treated is

$$\begin{aligned}
\sigma\left(\omega, \tau^{-1}_{\mathcal{G},\omega}(x)\right) &= \int_{1-x}^1 p^\omega(k)\,\mathrm{d}\omega_K \\
&= \int_{1-x}^1 k\mathrm{d}k \\
&= \tfrac{1}{2}\left(1-(1-x)^2\right) = \tfrac{1}{2}\left(2x - x^2\right)
\end{aligned}$$

From the decomposition $\sigma_{\mathcal{G},\omega}(x) = \phi_{\mathcal{G},\omega}(x)\,\sigma_\omega x$ it follows

[81] As usual, to have outcome a deterministic function of type we add a "reservation effort" characteristic $\overline{e} \sim \mathcal{U}[0,1]$, independent of k.

[82] Formally, $\overline{\tau}^{-1}_{\mathcal{G},\omega}(x)$ is a function from U to $[0,1]$ with $\overline{\tau}^{-1}_{\mathcal{G},\omega}(x)(u) = \mathbf{1}_{u\leq\widehat{u}}$ for some threshold $\widehat{u}$. I identify the function with the threshold.

$$\phi_{\mathcal{G},\omega}(x) = 2 - x$$

which is decreasing and always above 1: whatever the treated proportion is, realizing it through an incentivized regime proves more efficient than through coercion. This is intuitive, as we select firstly those agents that are willing also to put more effort. Consider the decomposition (2.24) and notice in this example $\sigma = \frac{1}{2}$, $\tau(q) = 1 - x$ and $\phi(x) = 2 - x$

$$\sigma'(q) = \underbrace{\frac{1}{2}}_{\sigma}\underbrace{(-1)}_{\tau'(q)}\left[\underbrace{2-(1-x)}_{\phi(\tau(q))} + \underbrace{1}_{\phi'(\tau(q))} \cdot \underbrace{1-x}_{\tau(q)}\right] = -\frac{1}{2}$$

Again recall we have an inverse order, so this means that increasing intensity marginally (decreasing price marginally) increases the success rate of $\frac{1}{2}\mathrm{d}q$. Notice that

$$\sigma(0) = \sigma(1) - \int_0^1 \sigma'(q)\,\mathrm{d}q = \frac{1}{2} = \sigma_\omega$$

The population outcome model

The empirical counterpart of the ATO is the number of citizens *in the sample* that have potential outcome equal to 1.

Definition 2.23. Given a sample ς, the **empirical success rate (ESR)** is $\overline{y}_\varsigma = |\{i \in \varsigma : y_i = 1\}|$.

The following proposition characterizes the joint behavior of EPS and ESR

Proposition 2.6. *It holds* $|\mathbb{E}_{\Sigma_N}[\overline{y}(g,\omega)] - \sigma(g,\omega)| \leq \frac{1}{N}$ *for all* N, ω, g. *Moreover, under the sampling procedure proposed in Section 2.5,*

$$\sqrt{N}\begin{pmatrix} \overline{\varsigma}_{g,\omega} \\ \overline{y}_{g,\omega} \end{pmatrix} \xrightarrow{d} \mathcal{N}\left(\begin{pmatrix} \tau(g,\omega) \\ \sigma(g,\omega) \end{pmatrix}, \begin{pmatrix} V^\tau_{g,\omega} & Cov^{\sigma,g}_{\tau,y} \\ \cdot & V^\sigma_{g,\omega} \end{pmatrix}\right)$$

for finite $V^\tau_{g,\omega}, V^\sigma_{g,\omega}$ *and* $Cov^{\sigma,g}_{\tau,y} = \sigma_{g,\omega}[1 - \tau(g,\omega)]$.

Finite-N properties of $\begin{pmatrix} \overline{\varsigma}_{g,\omega} \\ \overline{y}_{g,\omega} \end{pmatrix}$ are given in the proof. Notice that EPS and ERS are positively correlated because good outcome can result only from treatment: larger sample mechanically predicts higher ERS. It would be different if we had $\overline{y}$ the average treatment on the treated, that is $\frac{ESR}{EPS}$. The joint properties of those random variable would be more difficult to study, and beyond the scope of this paper.[83] In the proof it is given the finite sample joint distribution of those objects, which therefore gives the signal likelihoods in case no further (pre and post sample) characteristics were observed. In particular, notice that if the assignment mechanism is implicit (reservation price type does not enter planner's dataset), then results form the outcome phase provide additional information on the marginal ω_U that is not contained in observation of the EPS.

2.6.2 Unconfoundedness

Unconfoundedness is a fundamental assumption in the literature of estimation of treatment effect. Suppose a sample $\{Y, W, X\}$ is available, with Y an outcome, W the treatment status and X a set of (observable) covariates. Unconfoundedness is a statistical assumption on the available dataset

$$W \perp Y \big| X$$

That is, outcome is uninformative of selection choice given all the observable covariates. In the setting of this chapter, unconfoundedness is a joint property of the reduced form model and the selection mechanism (state and action in the MAB problem). We now give the super-population counterpart of this assumption[84]

Definition 2.24. A model and selection intensity pair ω, g are **jointly unconfounded under s**, written $\{g, \omega\} \in \upsilon_s$ if it holds[85]

$$g(\mathbf{u}) \perp^{\omega} \mathbf{y} \big| \mathbf{s}$$

[83] An intuitive property is that the two should not correlate in g, ω are unconfounded under the trivial characteristic.

[84] Notice that, the actual treatment realization is a random variable $\widehat{\boldsymbol{\tau}}$ parametrized by $g(u)$ only so that $\widehat{\boldsymbol{\tau}} \perp^{\omega, \Sigma_N} \mathbf{y} \big| \mathbf{s} \iff g(\mathbf{u}) \perp^{\omega} \mathbf{y} \big| \mathbf{s}$.

[85] This is the standard (conditional) statistical independence of two random variables defined on the same probability space.

A set of models $\Upsilon \subset \Omega$ and intensities $\mathcal{G} \subset I(U,[0,1])$ are **unconfounded under s** if

$$\forall g \in \mathcal{G}, \omega \in \Upsilon, \{g, \omega\} \in \upsilon_s$$

A set of models $\Upsilon \subset \Omega$ (a set of intensities $\mathcal{G}$) is **globally unconfounded under s** if it is unconfounded with $I(U,[0,1])$ (with Ω) under $\mathbf{s}$.

Unconfoundedness is a joint property of a reduced form model, a selection intensity and a characteristic. By definition, $\Omega, I(U,[0,1])$ are globally unconfounded under $\mathbf{u}$. The following Proposition characterizes properties of unconfounded pairs.

Proposition 2.7. *i) If $\omega, g \in \upsilon_s$, then*

$$\sigma_{g,\omega} = \int_S g(s)\, \sigma_\omega(s)\, \omega_S(\mathrm{d}s) \tag{2.25}$$

where $g(s) = \int_U g(u)\, \omega_{U|S}(\mathrm{d}u|s)$.

ii) A model ω is globally unconfounded under $\mathbf{s}$ *if and only if $\mathbf{u} \perp^\omega \mathbf{y} \big| \mathbf{s}$, that is $\sigma_\omega(\mathbf{u}, \mathbf{s}) : U \times S \to [0,1]$ is $\sigma(\mathbf{s})-$measurable.*

iii) A selection intensity g is globally unconfounded under $\mathbf{s}$ *if and only if for each $s \in S$, either y or $g(u)$ as functions from X to Ω^{RF} are constant on $s^{-1} := \{x \in X : s(x) = s\}$.*

iv) If $\{\mathcal{G}, \omega\} \in \upsilon_s$ and s is the trivial characteristic, then $\phi \equiv 1$ where ϕ is the regime distortion function in Definition 2.22.

Point $i)$ implies that under unconfoundedness we can write an alternative outcome model (2.25) in the characteristic that is unconfounded. For static optimality we would need to know the CATO $\sigma_\omega(s)$,[86] the *marginal* distribution of the characteristic ω_S, and the induced selection map $g(s)$. In the original success rate expression (2.23) we only had the CATO$-u$ and the marginal ω_U, as the induced selection intensity corresponded with the stimulus by minimal sufficiency of u. In the modified outcome model instead we must take into account that characteristic s *is not directly incentivizable* and that stimulus response for different values of the characteristic are determined by the correlation with the incentivizable statistic. Unconfoundedness

[86] Notice that the original outcome model needed the CATOs of the statistic sufficient for treatment $\sigma_\omega : U \to [0,1]$.

under observable covariates will make the estimates CATOs useful to inform about the success rate of alternative policies that discriminate incentivation across the covariate. If the covariate is observable only ex-post (effort in Example 2.9 below), information on the success rate of alternative policies is mediated by the induced selection function. What is relevant for the information acquisition decision of the planner is that if the belief is reasonably concentrated on set of models that satisfy unconfoundedness over a simple message,[87] it may be worth taking costly monitoring actions that induce disclosure of that message learn about the success rate of alternative policies.

Notice that if $\mathbf{y}$ is always nondegenerate conditional on $\mathbf{s}$ then condition *iii*) for a selection intensity to be globally unconfounded is equivalent to require[88]

$$\boldsymbol{u}^{-1}\left(U_g\right) \subseteq \sigma_X\left(\boldsymbol{s}\right)$$

where

$$\boldsymbol{u}^{-1}\left(U_g\right) := \sigma\left(\left\{x : u\left(x\right) \in A\right\}_{A \in U_g}\right)$$

and U_g is the identification domain defined in 2.16. It then follows immediately that if $U_g = \{\emptyset, U\}$, then g is globally unconfounded under any characteristic s. This would hold as an if and only if statement if s is the trivial characteristic.[89] Using this and Example 2.5 we get

Corollary 2.1. *The RCT regime is globally unconfounded under any characteristic. If* $\mathbf{y}$ *is non-degenerate, then only the RCT regime is globally unconfounded under the trivial characteristic.*

Finally we can use point *iv*) to rewrite the (2.24) for pairs unconfounded under the trivial characteristic as

$$\begin{aligned} \sigma'\left(q\right) &= \sigma\tau'\left(q\right)\left[\phi\left(\tau\left(q\right)\right) + \tau\left(q\right)\phi'\left(\tau\left(q\right)\right)\right] \\ &= \sigma\tau'\left(q\right)\left[1 + \tau\left(q\right)\cdot 0\right] \\ &= \sigma\tau'\left(q\right) \end{aligned}$$

Intensity impacts on the success rate only through its impact on the SP-PS, sample composition being irrelevant. As factorization (2.15) implies condition *ii*),

[87] In a static setup, unconfoundedness has no directly testable implications. However it becomes testable in this environment as we allow to observe dataset generate by the same reduced form model under different incentive schemes.

[88] As $\boldsymbol{u}^{-1}\left(U_g\right) \subseteq \sigma_X\left(\boldsymbol{s}\right)$ implies that $g\left(\mathbf{u}\right)$ is a function of $\mathbf{s}$.

[89] Hence $\sigma_X\left(\mathbf{s}\right) = \{\emptyset, X\}$

the vaccination example has models in the support Υ of planner's belief be globally unconfounded.[90] Hence,

$$\sigma'_{\omega}(q) = \mu_{\omega} \cdot \frac{\partial}{\partial q}\tau_{\omega}(q) = \mu_{\omega}\frac{\partial}{\partial q}\left(1 - e^{-\gamma q}\right) = \mu_{\omega}\gamma e^{-\gamma q}$$

and

$$\lim_{q\to\infty} \sigma_{\omega}(q) = \sigma_{\omega}(0) + \int_{\mathbb{R}_{+}} \mu_{\omega}\gamma e^{-\gamma q}\mathrm{d}q = \mu_{\omega}$$

which is consistent with μ_{ω} being the ATO.

In case $\mathcal{G}$ is the regime of RCTs, it further holds $\tau'(x) = 1$ and we get $\sigma_{\omega,g_{0,x}} = x\sigma_{\omega}$.

I conclude this subsection with an example (continuation of Example (2.2)) that clarifies that inside a decision model of treatment selection the unconfoundedness assumption under the trivial characteristics is essentially imposing impossibility of agent to affect outcome through post-treatment actions and that he does not have pre-selection private information.

Example 2.9. After being treated, citizens can choose effort $e \in \{0, 1\}$. They receive good outcome if and only if they get treated and the effort exerted exceeds the random reservation effort $\overline{\mathbf{e}}$ which is distributed with CDF F. Thus $\mathbf{y} = y(\overline{\mathbf{e}}, \mathbf{e}^*) = \mathbf{1}_{\overline{\mathbf{e}}\leq\mathbf{e}^*}$. Choice characteristics are $c = [r_0, r_1, u_1, k]$. The speculation structure has characteristic set $X = \mathbb{R}^5$, with $x = [\overline{e}, c]$. Agents are expected utility maximizer and the treatment-conditional problem reads

$$\max_{e\in\{0,1\}} U(c, e) = (r_0 + r_1 e)\, u_1 - k\mathbf{1}_{e=1}$$

(Hidden action) Define the random variables $[\mathbf{u}, \mathbf{e}^*] : X \to \mathbb{R}^2$ with $\mathbf{u} = V(\mathbf{x})$ the value of treatment (that is minimally sufficient for treatment under price incentivized assignment mechanisms) and $\mathbf{e}^* = e^*(\mathbf{x})$ the associated policy.

Consider the set of structural models $\Upsilon \subset \Omega$ characterized by

$$\omega \in \Upsilon \implies \omega = \omega_{\overline{E}} \times \omega_C \tag{2.26}$$

and

$$supp\,(\omega_{\overline{E}}) = \{a, b\}\,, \text{ with } a < 0 < 1 < b \tag{2.27}$$

[90]Notice however that the regime fails condition *iii*) and therefore would not be unconfounded under alternative models that correlate reservation price with outcome type.

that is, the outcome type is independent of decision types and is such that outcome does not depend on the hidden action.[91] It is immediate to check that Υ is totally unconfounded under the trivial characteristic.

Now consider an alternative set of models Υ', where only (2.26) holds. We have

i) Υ' *is not* totally unconfounded under the trivial characteristic; indeed the reservation value $\mathbf{u}$ will be correlated with the effort spent $\mathbf{e}^*$ conditional on treatment, and this drives the correlation even if the outcome type is distributed independently of all determinants of choice. However,

ii) Υ' *is* totally unconfounded under $\mathbf{s} = \mathbf{e}^*$: after conditioning for the effort spent, outcome is uninformative about the response to monetary stimuli in the treatment selection phase. Indeed, for each $\omega \in \Upsilon'$, $p^{\omega}(u, s) = F^{\omega}_{\overline{E}}(s)$ for $s \in \{0, 1\}$.

(Hidden information) The independence property (2.26) would fail in case individuals have private information about their reservation effort. Suppose agents observe a signal $s = \overline{e} + \nu\epsilon$, where ϵ is a standard normal random variable and $\nu \in \mathbb{R}_+$ is a measure of information accuracy. If we add the assumption that beliefs $r = [r_0, r_1]$ are derived from common prior, we are essentially imposing a correlation structure between $\overline{\mathbf{e}}$ and $\mathbf{r}$. Let $p^{\omega} \in \Delta\left(\overline{E}\right)$ be the common prior and $\left\{p^{\omega}_s \in \Delta\left(\overline{E}\right)\right\}_{s \in S}$ be the set of posteriors obtained after observing signal $s \in S$, P_s being the respective CDFs. Now, $supp\mathbf{r} = \{[P^{\omega}_s(0), P^{\omega}_s(1)]\}_{s \in \mathbb{R}}$ and the conditional density $f^{\omega}(\cdot\,|\overline{e}) : \overline{E} \to \mathbb{R}_+^{supp\mathbf{r}}$ is given by

$$f^{\omega}([P_s(0), P_s(1)]\,|\overline{e}) = \phi\left(\frac{s - \overline{e}}{\nu}\right) \tag{2.28}$$

Under no information ($\nu = \infty$) we are in a special case of (2.26);[92] under perfect information ($\nu = 0$) we have conditional belief is the vector $[\mathbf{1}_{\overline{e} \leq 0}, \mathbf{1}_{\overline{e} \leq 1}] \in \{0, 1\}^2$. If this is the case, agents know their health status conditional on each effort choice, and it follows that whenever they are willing to pay positive price they will have successful outcome, hence $\sigma(g, \omega) = \tau(g, \omega)$. Under intermediate information structures, (2.28) determines the correlation between beliefs and outcome types. Those beliefs impact both the reservation value and the treatment-conditional effort, and would therefore make unconfoundedness under $\mathbf{e}^*$ fail.

[91] $y(\overline{\mathbf{e}}, \mathbf{e}^*) = \mathbf{1}_{\overline{\mathbf{e}} = b}$. This formulation can be seen as a structural justification of the vaccination example.

[92] Special case, as the CPA implies that the support of the belief characteristics is a singleton.

2.6.3 The sequential treatment assignment problem

We now have all the elements to setup the problem of repeated price incentivation into social programs as a MAB problem described in Section 2. This is the aim of the remainder of this Section.

Control

At each period a set of outcome results as a function of planner's incentive scheme and the population drawn from the reduced form state and allocated to treatment and no treatment by a valid sampling procedure. This set of outcomes includes

- Number of treatments given, and the associated cost $q^* \cdot N\overline{\varsigma}$, where $N\overline{\varsigma}$ is the sample size.
- The enforcement revenue, resulting from different types paying for their treatment. It depends on the reservation price of citizens and the label associated to them.
- An empirical distribution over material outcomes $\Delta^{emp}(Y)$, $Y \in \{NT, 0, 1\}$. It is given by a three dimensional vector with integer entries that sum to N. In case $y_{NT} = 0$ we identify it with a single point, number of agents in good health status.

Planner ranks those outcomes according to a separable utility function $v : \mathbb{R}^2 \rightarrow \mathbb{R}$ given by

$$v(\Pi, \overline{y}) = \upsilon(\overline{y}) + \chi \Pi$$

where υ describes rankings over health outcomes, and Π are net profits for selling the treatment. To this utility is should be subtracted the direct experimentation cost from further monitoring actions that are taken. As in this subsection I am interested in static optimality, I disregard this last component as it will never be statically optimal to pay positive monitoring costs. Each realized sample gives information relevant for health outcome, but not for revenues (as it does not say the realized offer price in the BDM mechanism). However, separability and linearity in revenues allow to write

$$u(g, \omega) = \int \upsilon(\overline{y}(\varsigma))\, \mathrm{d}\mathbb{P}_{\Sigma_N(g,\omega)}(\varsigma) + \chi \int_U [t_g(u) - q^*]\, \mathrm{d}\omega_U(u) \tag{2.29}$$

where the measure $\mathbb{P}_{\Sigma_N(g,\omega)}$ was defined in (2.21), and

$$t_g(u) = \int_0^u u' \mathrm{d}g(u')$$

the integral on the RHS being a Stieltjies integral with respect to the increasing function g. Static optimality (disregarding information acquisition) is defined in the standard way

Definition 2.25. The **static policy correspondence** $G : \Omega^{RF} \rightrightarrows I(U, [0,1])$ is given by

$$G(\omega) = \arg\max u(g, \omega)$$

The following result characterizes optimality of the regime of conditionally deterministic selection intensities (see Example 2.4) under linear preferences in health outcomes.

Proposition 2.8. *If v is a linear function $v(\overline{y}) = k\overline{y}$, then $G(\omega) \cap CD \neq \emptyset$ for all ω, where CD is the regime of conditionally deterministic assignment mechanisms. In particular the optimal price is either in $\{0, \overline{u}\}$, or is a $\widehat{u}$ satisfying the condition (suppose ω_U admits density f^ω)*

$$[k\sigma_\omega(\widehat{u}) - \chi(u^* - \widehat{u})] f^\omega(\widehat{u}) = \chi(1 - F^\omega(\widehat{u})) \tag{2.30}$$

The result is a consequence of the problem with linear utility being a linear program with convex constraint set, so optima are extremum points of the constraint set. CD mechanisms (step functions) are extreme points in the space of increasing functions from U to the unit interval. Expression (2.30) makes clear that, in principle we need to know the whole outcome model, as the static policy depends on the density and outcome conditional on treatment for all possible reservation prices. Notice one can obtain the optimality condition for the vaccination example from (2.30), letting $\sigma_\omega(u) = \omega$ and substituting the exponential density. In case $\chi = 0$, then we would have $\widehat{u} = 0$: if planner cares only about the outcome then of course he will post reservation price null, having everyone enter the treatment. The case $\chi \to \infty$

has instead the usual profit maximization condition $f\left(\widehat{u}\right)=\frac{\left(1-F\left(\widehat{u}\right)\right)}{\widehat{u}-u^{*}}$ condition for a monopolist who knows the market demand.

If υ is linear, then so is u and certainty equivalence holds by having for each belief $\mu\in\Delta\left(\Omega^{RF}\right)$, $G\left(\mu\right)=G\left(\overline{\omega}_{\mu}\right)$ where

$$\overline{\omega}_{\mu}=\int_{\Omega^{RF}}\omega\mathrm{d}\mu\left(\omega\right)$$

static optimality implies choosing the CD mechanism that maximizes utility of the average reduced form model.

Proposition 2.8 has pretty strong hypothesis. Firstly, linearity of υ is a restrictive assumption (especially in cases the outcome space is more than two-dimensional). A popular utility function that fails linearity is the so-called voting criterion: the decision maker cares about re-election which occurs if and only if a percentage x of the population is satisfied (receives good outcome). In that case $\upsilon\left(\overline{y}\right)=\mathbf{1}_{\overline{y}\geq x}$. Consider the case in which $\upsilon\left(\overline{y}\right)$ depends on both $\mathbb{E}\left(\overline{y}\right)$ and $\mathbb{V}\left(\overline{y}\right)$. Under the proposed sampling scheme, RCT is implemented by making sure of the fraction that gets treated, while CD need to keep sampling stochasticity to preserve validity. Suppose the reduced form model has $p\left(u\right)\equiv\sigma$, and consider an RCT-$x$ and the associated CD with reservation $u_{x}:1-F\left(u_{x}\right)=x$ (same fraction x is treated expectation). Then $\overline{y}_{RCT-x}\sim Bin\left(\sigma,xN\right)$, while and in the second case $\overline{y}_{CD-u_{x}}\sim Bin\left(\sigma\left[1-F\left(u_{x}\right)\right],N\right)$ so that

$$\mathbb{V}\left(\overline{y}_{RCT-x}\right)=\frac{1}{N}\sigma\left(1-\sigma\right)x$$

$$\mathbb{V}\left(\overline{y}_{CD-u_{x}}\right)=\sigma\left[1-F\left(u_{x}\right)\right]\left[1-\sigma\left[1-F\left(u_{x}\right)\right]\right]\frac{1}{N}=\frac{1}{N}\sigma x\left[1-\sigma x\right]$$

The latter being greater, and planner needs to take into account the increased sampling uncertainty induced by the CD mechanism. If χ is small enough increased risk may overcome the revenue losses associate to a coercive mechanism.

Another concern with the hypothesis in Proposition 2.8 is that we disregard budget constraints in the form of a maximal quantity of treatment units available per period. CD mechanism give positive probability to any sample size, which would force the planner to default on its promise to provide treatment to anyone willing to accept the outcome of the BDM lottery. If budget constraint have to hold with probability one, then by point *iii*) of Proposition 2.5 the planner has to use a selection intensity bounded above by $\frac{M}{N}$, where M is the maximal units of treatment available.

Dataset

Beyond material outcomes, at the end of each period planner receives a set of signals (observations) from each citizen. The aggregate dataset is a function of the selected action and the realized population and sample.

$$D_t\left(P_t, \varsigma_t, a_t\right) = \left\{d^T_{a_t}\left(P_t\right), d^Y_{a_t}\left(\varsigma_t\right)\right\} \tag{2.31}$$

Set of signals contain messages sent by all citizens during the treatment phase and those sent during the outcome phase by those who end up in the sample ς_t. Notice the sample contains signals that would be sent under all possible actions, and the particular monitoring effort determines those that are actually observed. In particular examples, $d^T_{a_t}\left(P_t\right)$ would contain the reservation price if assignment is conducted through an explicit mechanism and nothing if it is implicit.[93] $d^Y_{a_t}\left(\varsigma_t\right)$ contains by assumption the resulting health outcome and it may contain the additional covariates that are chosen to be observed; a_t essentially determines the portion of signal characteristics of citizens in the sample that are (chosen to be) observed. Treatment status is implicitly observed as the sequence of health outcomes on the sample reveals the sample size. Now, let S_a the space of datasets that can realize after a is chosen. It includes U^N if a prescribes an explicit assignment mechanism, and apart from that it includes the space of random length samples of the messages that action a discloses. As an example, if only health outcomes are observed and the mechanism is implicit, then $S_a = \prod_{i=1}^N \{0,1\}^i$. If a' also monitors effort, then $S_{a'} = \prod_{i=1}^N \left(\{0,1\} \times E\right)^i$, if a'' has the assignment mechanism explicit then $S_{a''} = U^N \times \prod_{i=1}^N \left(\{0,1\} \times E\right)^i$. and so on.

Now fix a and pick $s \in S_a$. Let

$$S_a^{-1}\left(s\right) = \left\{P, \varsigma : D\left(P, \varsigma, a\right) = s\right\}$$

to write the likelihood function

$$L\left(s \middle| \omega, a\right) = \int_{S_a^{-1}(s)} \omega^N\left(\mathrm{d}P\right) \mathbb{P}_{\Sigma_N}\left(\mathrm{d}\varsigma \middle| P, g_a\right) \tag{2.32}$$

Those likelihoods then induced the state-action measures over the signal space $S = \bigcup S_a$. Thole likelihoods then determine the Bayes posterior and the Markov kernel maps that determine the property of a planner's action as a social experiment.

[93]By Remark, 2.1 which is a restatement of Proposition 7 in CPS, the assignment mechanism can elicit at most (i.e., if explicit) the determinants of selection choice, namely the reservation price.

The dynamic problem

We now have all the ingredients to set up the sequential incentive design problem. The action set $\mathcal{A} \subset I(U,[0,1]) \times \{0,1\}^S$ contains pairs of selection intensities and monitoring decisions (each 1 in the vector corresponds to monitoring a certain dimension). Planner chooses a stochastic process of selection intensities and monitoring pairs adapted to the filtration generated by information shocks that determine at each period the realized signal draw (2.32) from the distribution parametrized by a_t, ω. The problem is

$$\max_{a \in L(\Sigma, \mathcal{A}^\infty)} \mathbb{E}_{\mu,\nu^\infty} \sum_{t=0}^{\infty} \beta^t \left[u\left(g_{a_t}, \omega\right) - c\left(a_t\right)\right]$$

where g_a is the selection intensity prescribed by action a, $u\left(g_a, \omega\right)$ is given in expression (2.29), $c\left(a\right)$ are the monitoring costs associated to a that do not depend on the state, $\mu \in \Delta\left(\Omega\right)$ is the prior over reduced form models that results from a conjecturing exercise mediated by a speculation structure, and ν the measure over information shocks.

Using the recursive representation derived in Section 2 we can characterize, for I the DeGroot information function relative to the dynamic Bayes risk, the optimal assignment-monitoring action taken at belief state μ

$$a^*\left(\mu\right) \in \arg \max_{a \in I(U,[0,1]) \times \{0,1\}^S} u\left(g_a, \mu\right) + \left[I_\mu\left(a\right) - c\left(a\right)\right] \tag{2.33}$$

Both direct costs $c\left(a\right)$ and indirect costs from deviating from the static optimal incentive $g^* \in G\left(\mu\right)$ are traded off for increased information $I_\mu\left(a\right)$. Information comes by discretely changing the monitoring decisions to expand the post-treatment dataset, but also from changing the regime and intensity of the assignment mechanism.

2.7 Extensions And Conclusions

2.7.1 Non-coercive stimuli: Persuasion and awareness campaigns

We have focused our attention on monetary incentivation: each agent had to choose a treatment lottery from a set exogenously given to them. There are alternative ways in which a planner can alter individual selection choice.

Non-coercive stimuli have the planner take some actions that alter the individual evaluation of the treatment environment, thus influencing their selection choice without limiting their choice set. Two examples that may be relevant are persuasion and awareness campaigns. The former has the planner send a public message based on which agents update their beliefs about the subjective treatment environment (see subsection). Agents will respond differently based on their prior evaluation, preferences over outcomes and parameters that describe how much they trust planner's suggestion. A structural model will need to specify their super-population distribution. Given stationarity the problem is essentially a sequence of static multiple receiver Bayesian persuasion games with possibly heterogenous prior beliefs[94] and heterogeneous subjective updating rules. An awareness campaign provides instead agents with some technology to learn their outcome type; examples are free blood pressure tests that inform agents of their likelihood to suffer of hearth problems, orientation tests may reveal individual skills in the pursue of higher education. Agents are not persuaded through cheap talk, but provided with a technology that may help them discovering their outcome relevant type. Banerjee et al. (2007) and Weiss and Tschirhart (1994) have economic applications of information campaigns.[95]

Two concerns arise on whether persuasive incentive scheme can reasonably respect the assumptions made in Section 2.4. Firstly, stationarity is critical whenever individual beliefs are involved as they are not allowed to evolve over time. Here it is even more critical once we recognize that stationarity shuts down any reputation channel that is possibly relevant in a repeated persuasion game: the assumption is that an individual response to announcement would not change depending on the path of (verifiable) previous announcements. Awareness campaigns are more immune to this type of criticism. A second concern is that when agents take post-selection outcome relevant actions, campaigns that alter their perception of the environment they operate in may induce a different effort choice. As long as those choice alter the final outcome this class of stimuli will fail policy invariance that assumed individual outcome to be independent of the stimulus conditional on treatment choice. Reduced form model would need also outcome to depend on stimulus provided, or to study environments in which individual actions do not affect the likelihood of success (see the first part of Example 2.9).

Observable covariates

It is also reasonable to assume that some characteristics of the population are observed before choosing a stimulus. Observable characteristics partition the pop-

[94] The classic reference for Bayesian persuasion problems is Kamenica and Gentzkow (2011). Relevant extensions are, among others, Wang (2013) for multiple receivers and Alonso and Camara (2014) for heterogenous beliefs.

[95] Awareness campaigns are not prevention campaigns of the type studied in Kremer and Snyder Kremer and Snyder (2015): the latter are indeed conceptually closer to a stimulus inside this framework.

ulation into atoms of undistinguishable individuals. The reduced form model must be enlarged to describe the correlation between atoms: each atom will have its own reservation price - potential outcome - ex post observables joint distribution, and those distributions will be correlated across atoms. The latter correlation is a determinant of the external validity of experiments conducted on different atoms. The reduced form model answers now all external validity concerns (see later) by specifying model correlation across social groups. Clearly this comes at the cost of complicating the initial period conjecturing exercise.

Observable characteristics are essentially of two types, conceptually different:

- Demographics. Gender, race, family background.
- Social groups. The population may be naturally divided into groups. Student classes, workplaces, villages.

If for some (regulatory) reason incentivization cannot depend on observables, then we can proceed as in the previous sections but modify the dataset expression (2.31) enlarging $d_{a_t}(P_t)$ to include this covariate at no cost and irrespectively of a_t. With impossibility to condition incentivation, the distinction between ex-ante and ex-post observation is immaterial.[96] It may be relevant for estimation, however, as reduced form models may imply a correlation between such observables and characteristics that are outcome relevant.

More interesting is the case in which incentivization can be made contingent on observables. If this is the case, the incentivation scheme needs to prescribe a citizen-specific stimulus assignment subject to measurability with respect to the observable partition.[97] Ex-ante observables create a natural segmentation (see Bergemann et al. (2015)) in the market for treatment lotteries that can be exploited by the planner. This segmentation is important as it allows the planner to focus incentivization on classes that are more responsive or have higher reservation price. Different social groups are natural candidates to be experimentation units: we can provide an incentive scheme to evaluate its impact on a wider scale, or give different schemes to different villages (say, one RCT the other CD) and having the differential

[96]Still those characteristics would be different from a message sent during the outcome phase as they would enter the planner's information set irrespectively of the selection choice, this is why $d_{a_t}(P_t)$ and not $d_{a_t}(\varsigma_t)$ is modified.

[97]Clearly, with no observable the measurability condition reduces to everyone being subject to the same stimulus.

response in the two villages inform about the determinants of selection-outcome cross correlation. Experiments on social groups need to address the issue of external validity.[98] Banerjee et al. (2016) address external validity issues in a decision theoretic approach.

2.7.2 Conclusions

This paper studied dynamic decision problems with information acquisition in stationary environments. Section 2 derives for each sequential decision problem an associated measure of uncertainty in the sense of DeGroot (1962) that is decision relevant as that the determination of the optimal policy trades-off the implied information function one for one with expected utility at each belief state.

The econometric approach to estimation of treatment effect (Heckman (2008), Heckman and Vytlacil (2007)) is conceptually close to a (static) information acquisition problem. Stage 1 of the procedure specifies the decision problem $\mathcal{D}$ which incorporates the payoff relevant parameters, while the space of experiments $\mathcal{S}$ corresponds to the estimation procedures defined by functional form restrictions and statistical assumption on unobservables that the econometrician uses to make causal inference from an exogenously given dataset. Essentially, Blackwell signals are economic theories that discipline the estimation procedure and signal realizations are the estimates of parameter resulting from this procedure.[99] Chassang et al. (2012) point to the fact that institutional settings determine the set of experiments that can validly identify policy relevant parameters creates a potential to use incentive schemes as instruments for identification. If the incentive problem is repeated for several periods, each populated by random draws from an invariant distribution of individuals, policymaking becomes exploration in a bandit problem with correlated arms, where the returns are the superpopulation success rates (Definition 2.21) and the source of correlation is the reduced form model. Infinite repetition of the incentivation problem should be interpreted as a modeling simplification accepted by a planner that faces a "new" policy problem ($P3$ in Heckman and Vytlacil (2007), disease in the vaccination example) and recognizes that the accuracy with which he can estimate policy relevant parameters depend on the shape of the incentive scheme *and* the other

[98] $P2$ in Heckman and Vytlacil (2007): *forecasting the impacts of interventions implemented in one environment in other environments.*

[99] In this case there is not really a "cost" associated to each experiment if one does not take into account the complexity of the research effort required to the econometrician. The "best" signal is chosen, which adds a subjective component on the signal structure.

covariates that are ex-post observed, possibly as a result of a monitoring effort. The results from Section 2 provide the planner with an endogenous metric through which he can evaluate the informativeness of a selection intensity-monitoring effort pair in the same units of the *direct and indirect* costs associated to them.

References

Philippe Aghion, Patrick Bolton, Christopher Harris, and Bruno Jullien. Optimal learning by experimentation. *The review of economic studies*, 58(4):621–654, 1991.

Ricardo Alonso and Odilon Camara. Bayesian persuasion with heterogeneous priors. *Available at SSRN 2306820*, 2014.

Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.

Abhijit Banerjee, Rukmini Banerji, Esther Duflo, Rachel Glennerster, Daniel Kenniston, Stuti Khemani, and Marc Shotland. Can information campaigns raise awareness and local participation in primary education? *Economic and Political Weekly*, pages 1365–1372, 2007.

Abhijit Banerjee, Sylvain Chassang, and Erik Snowberg. Decision theoretic approaches to experiment design and external validity. Technical report, National Bureau of Economic Research, 2016.

Pierpaolo Battigalli, Mario Gilli, and M Cristina Molinari. *Learning and convergence to equilibrium in repeated strategic interactions: an introductory survey.* Università Commerciale" L. Bocconi", Istituto di Economia Politica, 1992.

Gordon M Becker, Morris H DeGroot, and Jacob Marschak. Measuring utility by a single-response sequential method. *Behavioral science*, 9(3):226–232, 1964.

Dirk Bergemann and Juuso Välimäki. Learning and strategic pricing. *Econometrica: Journal of the Econometric Society*, pages 1125–1149, 1996.

Dirk Bergemann and Juuso Valimaki. Bandit problems. 2006.

Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *The American Economic Review*, 105(3):921–957, 2015.

David A Blackwell. Comparison of experiments. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability*, volume 1, pages 93–102, 1951.

David A Blackwell and Meyer A Girshick. *Theory of games and statistical decisions.* Courier Corporation, 1979.

Patrick Bolton and Christopher Harris. Strategic experimentation. *Econometrica*, 67 (2):349–374, 1999.

Sébastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *arXiv preprint arXiv:1204.5721*, 2012.

Sylvain Chassang, Erik Snowberg, et al. Selective trials: A principal-agent approach to randomized controlled experiments. *The American Economic Review*, 102(4): 1279–1309, 2012.

Angus Deaton. Instruments, randomization, and learning about development. *Journal of economic literature*, 48(2):424–455, 2010.

Morris H DeGroot. Uncertainty, information, and sequential experiments. *The Annals of Mathematical Statistics*, pages 404–419, 1962.

David Easley and Nicholas M Kiefer. Controlling a stochastic process with unknown parameters. *Econometrica: Journal of the Econometric Society*, pages 1045–1064, 1988.

Pirmin Fessler and Maximilian Kasy. How to use economic theory to improve estimators, with an application to labor demand and wage inequality. 2016.

J Gittins and D Jones. A dynamic allocation index for the sequential allocation of experiments, in (j. gani, et al, eds.) progress in statistics, 1974.

John C Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 148–177, 1979.

David M Grether and Charles R Plott. Economic theory of choice and the preference reversal phenomenon. *The American Economic Review*, 69(4):623–638, 1979.

James J Heckman. Sample selection bias as a specification error (with an application to the estimation of labor supply functions), 1977.

James J Heckman. Econometric causality. *International statistical review*, 76(1):1–27, 2008.

James J Heckman and Edward J Vytlacil. Econometric evaluation of social programs, part i: Causal models, structural models and econometric policy evaluation. *Handbook of econometrics*, 6:4779–4874, 2007.

K Hinderer. Foundations of non-stationary dynamic programming with discrete time parameter. number 33 in lect. *Notes Op. Re. Math. Syst. Springer-Verlag, Berlin*, 1970.

Paul W Holland. Statistics and causal inference. *Journal of the American statistical Association*, 81(396):945–960, 1986.

Guido W Imbens and Donald B Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *The American Economic Review*, 101(6):2590–2615, 2011.

Edi Karni and Zvi Safra. " preference reversal" and the observability of preferences by experimental methods. *Econometrica: Journal of the Econometric Society*, pages 675–685, 1987.

Maximilian Kasy. Why experimenters should not randomize, and what they should do instead. *European Economic Association & Econometric Society*, 2013.

Ilan Kremer, Yishay Mansour, and Motty Perry. Implementing the "wisdom of the crowd". *Journal of Political Economy*, 122(5):988–1012, 2014.

Michael Kremer and Christopher Snyder. Preventives versus treatments. Technical report, National Bureau of Economic Research, 2015.

Sarah Lichtenstein and Paul Slovic. Reversals of preference between bids and choices in gambling decisions. *Journal of experimental psychology*, 89(1):46, 1971.

Yishay Mansour, Aleksandrs Slivkins, and Vasilis Syrgkanis. Bayesian incentive-compatible bandit exploration. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 565–582. ACM, 2015.

Yishay Mansour, Aleksandrs Slivkins, Vasilis Syrgkanis, and Zhiwei Steven Wu. Bayesian exploration: Incentivizing exploration in bayesian games. *arXiv preprint arXiv:1602.07570*, 2016.

Jacob Marschak. Economic measurements for policy and prediction. In *Economic information, decision, and prediction*, pages 293–322. Springer, 1953.

Giuseppe Moscarini. Job matching and the wage distribution. *Econometrica*, 73(2): 481–516, 2005.

Herbert Robbins. An empirical bayes approach to statistics. In *Herbert Robbins Selected Papers*, pages 41–47. Springer, 1985.

Michael Rothschild. A two-armed bandit theory of market pricing. *Journal of Economic Theory*, 9(2):185–202, 1974.

William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.

Abraham Wald. Foundations of a general theory of sequential decision functions. *Econometrica, Journal of the Econometric Society*, pages 279–313, 1947.

Abraham Wald. *Sequential analysis.* Courier Corporation, 1973.

Yun Wang. Bayesian persuasion with multiple receivers. *Available at SSRN 2625399*, 2013.

Janet A Weiss and Mary Tschirhart. Public information campaigns as policy instruments. *Journal of policy analysis and management*, 13(1):82–119, 1994.

Appendix

2.8 Proofs of Chapter 2

Proof of Proposition 2.3

For an increasing function g, the set Λ_g^{BDM} is well definite. I prove that it implements g. A lottery $L \in \Lambda_g^{BDM}$ takes the form

$$L = \langle [x, y_u], \mathrm{d}g(u) \rangle$$

where $x \in \{T, NT\}$ and $y_u \in \{0, -u\}$. By Definition 2.9 and money monotonicity MM, type u prefers $\langle [T, -u']; 1 \rangle$ to $\langle [NT, 0]; 1 \rangle$ if and only if $u' \leq u$. Then we can apply iteratively assumptions AAM (in its proper extension for distributions) to conclude

$$L^*(\Lambda_g, U)(u) = \{\mathcal{L}_{g,u'} : g(u') = g(u)\}$$

from which it follows $\mathcal{L}_{g,u} \in L^*(\Lambda_g, U)(u)$ the latter set being a singleton for us for which $g'(u) \neq 0$.[100] By construction $\tau(\mathcal{L}_{g,u}) = g(0) + \int_0^u \mathrm{d}g(u') = g(u)$, completing the proof of the if direction.

The only if part is proved by contrapositive assuming quasilinear preferences: individual of reservation price u ranks treatment probabilities and transfers according to $V(u, p, t) = pu - t$.[101] The IC constraint now takes the familiar form

$$u \in \arg\max_{u' \in U} g(u')\, u - t(u')$$

for all u. Assuming $u_1 > u_2$ and $g(u_1) < g(u_2)$ make the incentive compatibility constrains of u_1, u_2 yield the contradiction $t(u_1) < t(u_2) < t(u_1)$.

[100] By different from, it is included the case in which the object is not definite.
[101] It is clear that all assumptions are satisfied under this specification.

Proof of Proposition 2.1

Firstly I need to show that the Riemann integral on the RHS is well definite. Fix $g \in I(U,[0,1])$ and $\omega \in \Omega$. $U_g(x)$ is the pre-image through the function $g(u)$ of the set $[x,1] \in \mathcal{B}([0,1])$. Measurability of that function guarantees that for all $x \in [0,1]$, $U_g(x) \in \mathcal{B}(U)$, so that $\omega(U_g(x))$ is well defined self-map on $[0,1]$. As g is weakly increasing, $(U_g(x))_{x\in[0,1]}$ is a weakly decreasing (in the inclusion order) sequence of sets, then $\omega_U(U_g(x))$ is a bounded monotonic (decreasing) function on a compact set, hence it is Riemann integrable and the RHS of (2.18) is well definite.

Then we can write

$$\tau(g,\omega) = \int g(u)\,\mathrm{d}\omega_U(u) = \int_0^1 \omega_U(\{u : g(u) \geq x\})\,\mathrm{d}x = \int_0^1 \omega_U(U_g(x))\,\mathrm{d}x$$

The first and last equality just apply definitions given in the paper, while the second is the defining property of a Lebesgue integral.

Proof of Proposition 2.4

Take g a generic assignment function. For all $u \in U$

$$\begin{aligned} \mathbb{P}_{\Sigma_N}[\widehat{\tau}_g(u,\mathbf{k})=1] &= \mathbb{E}_{\Sigma_N}\left[\mathbf{1}_{\left[g(u)\geq\frac{\mathbf{k}}{N}\right]}\right] \\ &= \tfrac{1}{N}\left|\left\{k\in[N] : \tfrac{k}{N}\leq g(u)\right\}\right| \\ &= \lfloor g(u)\rfloor_N \end{aligned} \tag{2.34}$$

where in the second line $|\cdot|$ is the cardinality function and in the third line

$$\lfloor x\rfloor_N \coloneqq \max_{j\in\{0\}\cup[N]} : \frac{j}{N} \leq x \tag{2.35}$$

Now $i)$ follows from $\lfloor x\rfloor_N = x$ for $x \in \{0,1\}$.[102] As for $ii)$, individualistic assignment is obvious as, conditional on individual reservation type u, the sampling procedure depends only on type independent label assignment. The validity condition follows by plugging the expression (2.34) into (2.20) and noting $\left|\lfloor x\rfloor_N - x\right| \leq \frac{1}{N}$ for all x.

Proof of Proposition 2.5

$i)$ The procedure has each label being assigned randomly and independently to a citizen, while different labels have different type-conditional treatment assignment.

[102]In particular, for any $x \in \left\{\frac{j}{N}\right\}_{j\in\{0\}\cup[N]}$.

So to each label a reservation price type is assigned which is an independent draw of $\mathbf{u}$. By independence across labels it holds

$$\begin{aligned}\mathbb{E}_{\Sigma_N}\left[\overline{\varsigma}\left(g,\omega\right)\right] &= \tfrac{1}{N}\textstyle\sum_{k=1}^{N}\mathbb{E}_{\omega}\left[\widehat{\tau}_g\left(\mathbf{u},k\right)\right] \\ &= \tfrac{1}{N}\textstyle\sum_{k=1}^{N}\mathbb{E}_{\omega}\left[\mathbf{1}_{\left[g(\mathbf{u})\geq\frac{k}{N}\right]}\right] \\ &= \tfrac{1}{N}\textstyle\sum_{k=1}^{N}\mathbb{E}_{\omega}\left[\mathbf{1}_{\left[\mathbf{u}\geq g^{-1}\left(\frac{k}{N}\right)\right]}\right] \\ &= \tfrac{1}{N}\textstyle\sum_{k=1}^{N}\omega\left(U_g\left(\frac{k}{N}\right)\right)\end{aligned} \tag{2.36}$$

where g^{-1} and $U_g(x)$ used in the second and last inequality were defined in Definition 2.14. As each realization , the expectation completely describes the distribution and we have

$$\overline{\varsigma}\left(g,\omega\right)\sim_{\Sigma_N}\frac{1}{N}\sum_{k=1}^{N}Ber\left(\omega_U\left(U_g\left(\frac{k}{N}\right)\right)\right) \tag{2.37}$$

Which is a Poisson-Binomial random variable with parameters $\left[\left\{\omega_U\left(U_g\left(\frac{k}{N}\right)\right)\right\}_{k\in[N]}\right]$. Point i) then follows from

$$\left(\omega_n,g_n\right)\to\left(\omega,g\right)\implies\left[\left\{\omega_{U,n}\left(U_{g_n}\left(\frac{k}{N}\right)\right)\right\}_{k\in[N]}\right]\to\left[\left\{\omega_U\left(U_g\left(\frac{k}{N}\right)\right)\right\}_{k\in[N]}\right]$$

and the Poisson-Binomial density being continuous in its parameters.

ii) Recall that, by Lemma 2.1 $\tau\left(g,\omega\right)=\int\omega_U\left(U_g\left(x\right)\right)\mathrm{d}x$ and that $\omega_U\left(U_g\left(x\right)\right)$ is a monotone decreasing function of x. Now we observe that (2.36) is just the Riemann approximation of the (Lebesgue) integral defining $\tau\left(g,\omega\right)$, evaluating the function at the endpoints of the equipartitions. Since the function is monotonically decreasing we have $\mathbb{E}_{\Sigma_N}\left[\overline{\varsigma}\left(g,\omega\right)\right]\leq\tau\left(g,\omega\right)$, and

$$\tau\left(g,\omega\right)-\mathbb{E}_{\Sigma_N}\left[\overline{\varsigma}\left(g,\omega\right)\right]\leq\frac{1}{N}\left(\omega_U\left(U_g\left(\frac{1}{N}\right)\right)-\omega_U\left(U_g\left(1\right)\right)\right)\leq\frac{1}{N}$$

which proves the first part of ii). The second part follows by the strong law of large numbers applied to a sequence of independent (but non identically distributed) random variables with uniformly bounded variance.

Under the assumptions of iii) it holds

$$U_g\left(\frac{k}{N}\right)=\begin{cases}U & \forall k\leq N\left\lfloor a\right\rfloor_N\\ \emptyset & \forall k>N\left\lfloor b\right\rfloor_N\end{cases},$$

So that, the first $N\lfloor a\rfloor_N$ parameters of the Poisson-Binomial are $\omega_U(U)=1$ irrespectively of ω, and similarly the last $N-N\lfloor b\rfloor_N$ are 0 irrespectively of ω. Hence,

$$\begin{array}{rcl}\overline{\varsigma}(g,\omega) & \sim & \frac{1}{N}\left(\sum_{k=1}^{N\lfloor a\rfloor_N} Ber(1)+\sum_{k=N\lfloor b\rfloor_N+1}^{N} Ber(0)+\sum_{k=N\lfloor a\rfloor_N+1}^{N\lfloor b\rfloor_N} Ber\left(\omega_U\left(U_g\left(\frac{k}{N}\right)\right)\right)\right) \\ & \sim & \lfloor a\rfloor_N+\frac{1}{N}\left(\sum_{k=N\lfloor a\rfloor_N+1}^{N\lfloor b\rfloor_N} Ber\left(\omega\left(U_g\left(\frac{k}{N}\right)\right)\right)\right)\end{array}$$

From which $iii)$ follows.

$iv)$ is a special case of $iii)$ when $a=b=x$.

Finally, notice that under a conditionally deterministic regime the post-draw type is irrelevant, that is for all $k\in[N]$

$$U_{\infty,u}\left(\frac{k}{N}\right)=U_{\infty,u}(1)=[u,\overline{u}]$$

independent of k. Then (2.37) implies that $\overline{\varsigma}(g,\omega)$ is the sum of independent Bernoulli random variables with parameter $\omega_U([u,\overline{u}])$ and $v)$ follows from $\tau(g_{\infty,u},\omega)=\omega_U([u,\overline{u}])$ by Example 2.3.

Proof of Proposition 2.7

$i)$ It holds

$$\begin{array}{rcl}\sigma_{g,\omega} & = & \int \sigma_\omega(u)\,g(u)\,\omega_U(\mathrm{d}u) \\ & = & \int_U\left(\int_S \sigma_\omega(s,u)\,\omega_{S|U}(\mathrm{d}s|u)\right)g(u)\,\omega_U(\mathrm{d}u) \\ & = & \int_U\left(\int_S \sigma_\omega(s)\,\omega_{S|U}(\mathrm{d}s|u)\right)g(u)\,\omega_U(\mathrm{d}u) \\ & = & \int_{U\times S}\sigma_\omega(s)\,g(u)\,\omega_{U\times S}(\mathrm{d}u,\mathrm{d}s) \\ & = & \int_S\sigma_\omega(s)\left(\int_U g(u)\,\omega_{U|S}(\mathrm{d}u|s)\right)\omega_S(\mathrm{d}s)\end{array}$$

Which is our desideratum. Unconfoundedness assumption was used in the third equality to have $p^\omega(s,u)=\sigma_\omega(s)$ (see point $ii)$ below). Other equalities just apply definitions and Fubini's theorem.

$ii)$ The if direction is immediate. Conversely, suppose $\mathbf{u}\not\perp^\omega\mathbf{y}|\mathbf{s}$. To show failure of total unconfoundedness consider a strictly increasing selection intensity g, and notice

$$\omega\left(\mathbf{u}\in A,\mathbf{y}\in B|\sigma(s)\right)\neq\omega\left(\mathbf{u}\in A|\sigma(s)\right)\cdot\omega\left(\mathbf{y}\in B|\sigma(s)\right)\implies$$

$$\omega\left(g\left(\mathbf{u}\right)\in g\left(A\right),\mathbf{y}\in B\middle|\sigma\left(s\right)\right)\neq\omega\left(g\left(\mathbf{u}\right)\in g\left(A\right)\middle|\sigma\left(s\right)\right)\omega\left(\mathbf{y}\in B\middle|\sigma\left(s\right)\right)$$

the first line holds by contrapositive hypothesis for some $A\in\mathcal{B}\left(U\right)$, $B\in\mathcal{B}\left(\{0,1\}\right)$, then it is used $g:U\to[0,1]$ being invertible (as g is strictly increasing) to equate both sides of the first inequality and obtain the second inequality. As $g\left(A\right)\in\mathcal{B}\left([0,1]\right)$ it follows $g\left(\mathbf{u}\right)\not\perp^{\omega}\mathbf{y}\middle|\mathbf{s}$.

iii) The if direction is trivial as any random variable is independent of a degenerate random variable. Conversely, suppose there is $s\in S$ and $\{x_i\}_{i=1}^{4}\in s^{-1}$ such that $g_1\neq g_2$ and $y_3\neq y_4$.[103] Notice that x_1 and x_2 may coincide with x_3 and x_4 (but clearly $x_1\neq x_2$). Consider a model $\omega\in\Omega$ such that $supp\left(\omega\right)\cap s^{-1}=\{x_i\}_{i=1}^{4}$ and $\omega\left(x_i\right)\neq\omega\left(x_j\right)$ $\forall i\neq j\in[4]$. As

$$\omega\left(y_i\middle|g_j,s\right)=\frac{\mathbf{1}_{y_i=y_j}\omega_i}{\sum_{i\in[4]}\mathbf{1}_{g_i=g_j}\omega_i}$$

It is a simple exercise to check that $g\left(\mathbf{u}\right)\not\perp^{\omega}\mathbf{y}\middle|\mathbf{s}=s$.

iv) As s is the trivial characteristic, $\sigma_\omega\left(s\right)=\sigma_\omega$ and $\omega_{U|S}\left(\mathrm{d}u\middle|s\right)=\omega_U\left(\mathrm{d}u\right)$. Now we use *i*) to write

$$\phi_{\mathcal{G},\omega}\left(x\right)=\frac{\int_S\tau_{\mathcal{G},\omega}^{-1}\left(x\right)\left(s\right)\sigma_\omega\left(s\right)\omega_S\left(\mathrm{d}s\right)}{\sigma_\omega x}=\frac{\int_U\tau_{\mathcal{G},\omega}^{-1}\left(x\right)\left(u\right)\sigma_\omega\omega_U\left(\mathrm{d}u\right)}{\sigma_\omega x}=1$$

where in the last equality we used

$$\int_U\tau_{\mathcal{G},\omega}^{-1}\left(x\right)\left(u\right)\omega_U\left(\mathrm{d}u\right)=x$$

by definition.

Proof of Lemma 2.6

Following the same steps as in the proof of Proposition 2.5 we have

$$\begin{aligned}\mathbb{E}_{\Sigma_N}\left[\overline{y}\left(\omega,g\right)\right] &= \frac{1}{N}\sum_{k=1}^{N}\mathbb{E}_\omega\left[\widehat{\tau}_g\left(\mathbf{u},k\right)\cdot\mathbf{y}\right]\\ &= \frac{1}{N}\sum_{k=1}^{N}\mathbb{E}_\omega\left[\mathbf{1}_{\left[g(\mathbf{u})\geq\frac{k}{N}\right]}\mathbf{y}\right]\\ &= \frac{1}{N}\sum_{k=1}^{N}\mathbb{E}_\omega\left[\mathbf{1}_{\left[\mathbf{u}\geq g^{-1}\left(\frac{k}{N}\right)\right]}\right]\cdot\mathbb{E}_\omega\left(\mathbf{y}\middle|\mathbf{u}\geq g^{-1}\left(\frac{k}{N}\right)\right)\\ &= \frac{1}{N}\sum_{k=1}^{N}\omega_U\left(U_g\left(\frac{k}{N}\right)\right)\mathbb{E}_\omega\left(\mathbf{y}\middle|\mathbf{u}\geq g^{-1}\left(\frac{k}{N}\right)\right)\end{aligned}$$

[103] Where g_i,y_i are shorthand to write $g\left(u\left(x_i\right)\right),y\left(x_i\right)$ respectively.

So that

$$\overline{y}(\omega, g) \sim PoissBin\left[\left\{\omega_U\left(U_g\left(\frac{k}{N}\right)\right)\cdot\mathbb{E}_\omega\left(\mathbf{y}\middle|\mathbf{u}\geq g^{-1}\left(\frac{k}{N}\right)\right)\right\}_{k\in[N]}\right]$$

Now define the function $\sigma_{g,\omega}:[0,1]\to[0,1]$ by

$$\sigma_{g,\omega}(x) = \left\{\left[\omega_U\left(U_g(x)\right)\right]\cdot\mathbb{E}_\omega\left(\mathbf{y}\middle|\mathbf{u}\geq g^{-1}(x)\right)\right\}$$

So that

$$\mathbb{E}\left[\overline{y}_{g,\omega}\right] = \tfrac{1}{N}\textstyle\sum_{k=1}^{N}\sigma_{g,\omega}\left(\tfrac{k}{N}\right)$$

Using the same argument as in Lemma 2.1 we can the write

$$\sigma_{g,\omega} = \int_0^1 \sigma_{g,\omega}(x)\,\mathrm{d}x$$

and the result follows from from the same approximating arguments as in Proposition 2.5 *ii*).

By independence across labels,

$$Var\left(\overline{\varsigma}_{g,\omega}\right) = \frac{1}{N}\sum_{k=1}^{N}\omega\left(U_g\left(\frac{k}{N}\right)\right)\left[1-\omega\left(U_g\left(\frac{k}{N}\right)\right)\right]$$

$$Var\left(\overline{y}_{g,\omega}\right) = \frac{1}{N}\sum_{k=1}^{N}\sigma_{g,\omega}\left(\frac{k}{N}\right)\left[1-\sigma_{g,\omega}\left(\frac{k}{N}\right)\right]$$

So that

$$Var\left(\overline{\varsigma}_{g,\omega}\right) \to V_{g,\omega}^{\tau} = \int \omega_U\left(U_g(x)\right)\left[1-\omega_U\left(U_g(x)\right)\right]\mathrm{d}x$$

$$Var\left(\overline{y}_{g,\omega}\right) \to V_{g,\omega}^{\sigma} = \int \sigma_{g,\omega}(x)\left[1-\sigma_{g,\omega}(x)\right]\mathrm{d}x$$

that are both finite. Finally, using label independence once again we can write

$$\begin{array}{rcl} Cov\left(\overline{\varsigma}_{g,\omega},\overline{y}_{g,\omega}\right) & = & \frac{1}{N}\sum_{k=1}^{N} Cov_{\omega}\left(\widehat{y}\left(\frac{k}{N}\right),\widehat{\tau}_{g}\left(\frac{k}{N}\right)\right) \\ & = & \frac{1}{N}\sum_{k=1}^{N}\mathbb{E}_{\omega}\left(\widehat{y}\left(\frac{k}{N}\right)\cdot\widehat{\tau}_{g}\left(\frac{k}{N}\right)\right)-\mathbb{E}_{\omega}\left(\widehat{y}\left(\frac{k}{N}\right)\right)\mathbb{E}_{\omega}\left(\widehat{\tau}\left(\frac{k}{N}\right)\right) \\ & = & \frac{1}{N}\sum_{k=1}^{N}\mathbb{E}_{\omega}\left(\widehat{y}\left(\frac{k}{N}\right)\right)\left[1-\mathbb{E}_{\omega}\left(\widehat{\tau}\left(\frac{k}{N}\right)\right)\right] \\ & = & \frac{1}{N}\sum_{k=1}^{N}\sigma_{g,\omega}\left(\frac{k}{N}\right)\left[1-\omega_{U}\left(U_{g}\left(\frac{k}{N}\right)\right)\right] \end{array}$$

where the third line uses $\widehat{y}\left(\frac{k}{N}\right)=1 \implies \widehat{\tau}_{g}\left(\frac{k}{N}\right)=1$ as only those who are treated can get good outcome. Now it is immediate to pass to the limit, use the result proved above and Proposition 2.5 *i*) to conclude

$$Cov\left(\overline{\varsigma}_{g,\omega},\overline{y}_{g,\omega}\right)\rightarrow\sigma_{g,\omega}\left[1-\tau\left(g,\omega\right)\right]$$

And the asymptotic result follows from application of the Lyapunov CLT.[104]

Proof of Proposition 2.8

If υ is linear, expression (2.29) simplifies to

$$\begin{array}{rl} & k\mathbb{E}_{\Sigma_{N}(g,\omega)}\left(\overline{y}\left(\varsigma\right)\right)+\chi\int_{U}\left[t_{g}\left(u\right)-q^{*}\right]\mathrm{d}\omega_{U}\left(u\right) \\ = & k\sigma\left(g,\omega\right)+\chi\int_{U}\left[t_{g}\left(u\right)-q^{*}\right]\mathrm{d}\omega_{U}\left(u\right) \end{array} \quad (2.38)$$

The proposition is proved for U with finite support. Extension is immediate due to density of simple functions. With finite support of cardinality m, the outcome model is a pair $f,p\in\mathbb{R}^{m}$, the former giving the mass at each point in the support, the latter the reservation price conditional likelihood of good health outcome. The discrete-support version of equation (2.38) is

$$\sum_{i=1}^{m}\left[kp_{i}-\chi q^{*}\right]f_{i}g_{i}-\chi\sum_{i=1}^{m}f_{i}\underbrace{\sum_{j\leq i}g_{j}u_{j}}_{t_{g}(u_{i})}$$

so that the discrete version of the maximization problem can be written as

$$\max_{g\in M}\sum_{i=1}^{m}\gamma_{i}g_{i}$$

where

$$\gamma_{i}=\left[\left(kp_{i}-\chi q^{*}\right)f_{i}+\chi u_{i}\sum_{j\geq i}f_{j}\right]$$

[104] Actually, this is not immediate as the stochastic process changes as we change n.

and M is the simplex over $\mathbb{R}^m$, namely

$$M = \{x \in \mathbb{R}^m : x \geq \mathbf{0} \text{ and } x \cdot \mathbf{1} \leq 1\}$$

This is a linear program under convex constraint set, which has solution extremum points. Extrema are vectors with entries $0, 1$ and (at most) a one. Those are exactly the discrete counterpart of CD mechanism, that assign unit mass to a single point.

Now we know that the optimal solution to the problem

$$\max_{g \in I(U,[0,1])} \int_U k\sigma(u) f(u) g(u) \, \mathrm{d}u - \chi \int_U \left(q^* g(u) - \left(\int_0^u u' \mathrm{d}g(u') \right) \right) f(u) \, \mathrm{d}u$$

takes the simpler form

$$\max_u \int_u^{\overline{u}} \left[k\sigma(u') - \chi(q^* - u) \right] f(u') \, \mathrm{d}u'$$

First order condition reads

$$\left[k\sigma(\widehat{u}) - \chi(q^* - \widehat{u}) \right] f(\widehat{u}) + \chi \int_{\widehat{u}}^{\overline{u}} f(u') \, \mathrm{d}u' = 0$$

which can be rearranged to obtain (2.30). This characterizes necessary conditions for an interior optimum. Global optimum compares the value of the function at those critical points with the boundary $\{0, \overline{u}\}$, from which the result follows.

Chapter 3

Screening for Susceptibility and Influence

(With Franz Ostrizek)

3.1 Introduction

Complementarities are a key feature of many markets. Consumers choose software and services like Dropbox and Facebook taking into account how many of their acquaintances will use them as well and base their purchase decisions on the advice of social media influencers in product categories ranging from makeup to consumer electronics. Investors rely on other firms for inputs and services and in turn provide them to other firms, creating complementarities.[1] Not all agents are equally dependent on inputs and content created by others or susceptible to peer pressure and fashion. Similarly, not all agents are equally influential among their peers or provide essential services to other firms. This heterogeneity creates incentives for the firm selling a network good or a planner taxing an industry to discriminate based on *susceptibility* and *influence*: Sell at high prices to consumers that are susceptible in order to exploit their high willingness to pay due to complementarities and sell large amounts cheaply to influential consumers in order to increase the willingness to pay of the population.

Reliable information on these characteristics can be hard to come by. Outside of core social networking services like Facebook and Twitter, the susceptibility and influence of individuals is rarely observable. Even in those applications, these measures such as follower counts can be easily manipulated by buying followers or

[1]See Bryant (1983); Cooper and John (1988); Benhabib and Farmer (1994); Ball and Romer (1990) and the following literature for the macro impact of investment and price setting complementarities.

creating spurious links (following others only to get discounts). Hence, markets with consumption externalities naturally create a screening problem with two dimensions of heterogeneity as firms attempt to price based on susceptibility and influence of their consumers.

In this paper, we build on a model in which complementarities are characterized by a weighted average of agents' consumption (from now on referred to as *aggregate consumption*), where the weights are proportional to influence. We study how the observability of susceptibility and influence affects the monopoly allocation. Clearly, if both characteristics are observed by the monopolist, she can implement the efficient contract and extract all rents. We show that this result holds as long as susceptibility is observed.

We then turn to the screening problem where both characteristics are unobserved. We exploit the special structure of the resulting two-dimensional screening problem with externalities to arrive at a tractable solution. As in a standard screening problem, incentive compatibility requires that the allocation is increasing in the level of susceptibility to the network effect. For a given level of susceptibility, the monopolist "tilts" the allocation to provide higher consumption to more influential consumers in order to provide a larger network externality. The two dimensions are treated differently, because the level of influence of a consumer doesn't directly enter his utility function. We show that the allocation increases in the lexicographic order, where susceptibility is the dominant dimension. We can hence transform the problem into a one-dimensional problem along this order. We solve it using standard techniques, that need to be generalized to take into account that the network effect creates interdependence among individuals' virtual values.

The virtual value will typically be non-monotonic in the lexicographic order for two reasons, both arising around the switching types (types with the highest level of influence, that are consequently adjacent to a type with higher susceptibility in the lexicographic order): First, only the consumption of these types directly causes information rents and hence only their virtual value is downward distorted. Second, the subsequent type in the lexicographic order has much lower influence. As influence positively enters the virtual value, this downward jump in influence can create non-monotonicity. So, given the additional concerns for non-monotonicity, the optimal contract will typically entail ironing of the allocation, and every bunching region will include at least one switching type.

One feature of the optimal contract with unobservable susceptibility and influence is that individuals' influence is not rewarded. Influential consumers receive the same

	susceptibility observable	susceptibility not observable
influence observable	Section 3.3.1	Section 3.4
influence not observable	Section 3.3.2	Section 3.5

Table 3.1: The four different observability assumptions.

level of utility as their less influential peers but consume more at lower unit price. This parallels the result for observable susceptibility. The drivers of the results are, however, different: In the latter case full surplus extraction left everyone without any rent, while in the former case incentive compatibility prevents any rent to emerge along a dimension (influence) that does not directly affect individuals' utility.

We also study the problem with observable influence (but unobserved susceptibility). In this case, a condition on primitives ensures that the optimal contract exhibits influence rents.[2] This shows that the economic gain of an influencer from his position crucially depends on its verifiability.

Natural restrictions on the possible misreports by consumers map into the observability assumptions we analyze exhaustively: When consumers can only over-report their types, the solution will be efficient, as if susceptibility were observable. When consumers can only under-report their types we find a condition under which the solution is the same as if only influence were observable.

This chapter proceeds as follows. We conclude this introductory section by discussing the relevant literature. Section 3.2 presents the model, which is conveniently characterized by only two equations: a function mapping individual consumptions into the aggregate consumption (first equation), which in turns affects the returns to individuals' consumption (second equation). We detail the application to non-linear pricing of a network good, and also discuss other economic problems that naturally fit into our framework once influence and susceptibility are appropriately redefined. The remaining sections focus on characterizing the revenue maximizing contracts under all four assumptions on observability of susceptibility and influence, as outlined in Table 3.1. Depending on the application, the principals' objective may depend on the consumption aggregator beyond revenues. Solving for the optimal (marginal) revenue associated to each aggregate consumption level (the focus of this paper) is however instrumental in finding the optimal contract with the augmented objective

[2] In particular, the condition puts an upper bound on the affiliation between influence and susceptibility.

function.[3] Section 6 discusses how restrictions on the possible misreports (over- and under-reporting) map into observability assumptions. Those results are important because – again depending on the particular application – restrictions on misreport emerge naturally as "technological constraints". Mapping those restrictions into observability assumptions which deliver a simple solution significantly simplifies the analysis of outcomes we expect would arise in those markets. Section 7 concludes. Most of the proofs can be found in the Proof Appendix.

3.1.1 Literature

Our model relates to the classic literature on contracting with network effects (Segal, 1999). Jadbabaie and Kakhbod (2016) compare bilateral and multilateral contracting in this setting when there are finitely many consumers and consequently, there is aggregate uncertainty about the realized distribution of types. This literature focuses on the externality of contracting in a setting with finitely many agents, in particular on the effect on the outside option of the agents. while we focus on a continuum of consumers with public contracting in a setting where the outside option of the agents is independent of the contract accepted by others. Sundararajan (2004); Csorba (2008) study screening with externalities in consumption when consumers have private information about their valuation of the good. We study screening on susceptibility and influence to the externality.

Our main application relates to the classic literature on externalities in consumption following the seminal Farrell and Saloner (1985) and Katz and Shapiro (1985). A recent literature focuses on the use of network information by a monopolist, both in the case of an explicit finite network (e.g. Bloch and Quérou, 2013; Candogan et al., 2012) and when consumers only know their level of susceptibility and influence (Fainmesser and Galeotti, 2016a,b). We adopt the demand and interaction specification developed in the latter, but focus on the screening problem.

There is a growing literature on monopolist screening for these characteristics. Zhang and Chen (2017) consider an explicit stochastic network formation model, where the out-degree of agents is fixed and consider screening along the in-degree. They consider two specifications, susceptibility is either a consumers in- or out-degree. Depending on this choice, their model can generate both quantity discounts and premia. Gramstad (2016) consider screening in a undirected network when network effects only depend on the number of neighbors that adopted the good, not their

[3] Actually, a characterization of marginal revenues is sufficient for most of the comparative statics we may be interested in the extensions tied to the particular applications where the additional benefit is non-zero.

intensity of consumption. We analyze both dimensions of private information – susceptibility and influence – at the same time and study their interaction in screening. The paper closest to ours is Shi and Xing (2018). They study screening with the same demand specification, but assume a continuum type space. As a result, the optimal allocation is constant in influence and the solution is one-dimensional. They focus on the implications for the value of network information, while we focus on the dependence of the allocation and rents on influence.

3.2 Model and Relevant Applications

This section presents the basic equations of the model and discusses the relevant applications. We construct a parsimonious model of two dimensional screening with externalities where the primitives are *i*) the map from individual consumptions into the aggregate (network) effect, which enters directly *ii*), individual utilities and *iii*) the principal's objective function. Heterogeneity along the influence dimension will affect only *i*), the impact of individual consumption in the creation of the externality, while susceptibility will be a parameter of the utility function *ii*), thus affecting the individual's sensitivity to the network effect. Although the main motivation of our paper - as highlighted in the introductory section - is the study of optimal non-linear pricing of a network good, we present three other economic problems that naturally fit the parsimonious description offered by primitives *i*) and *ii*) and to which our results can therefore be applied: non-linear taxation of externality producing activities, a price setting game among firms selling substitute goods, and subsidization of different sectors whose demands are shifted by a common aggregate measure of quality (big push development).

3.2.1 Setup and Primitives

There is a unit mass of agents consuming a divisible good x. The use of the good is subject to network (or aggregate) effects: the attractiveness of the good is dependent on the aggregate level of consumption $\bar{x}$. Each consumer is characterized by her level of susceptibility (to network effects) $k \in \mathcal{K} = \{\underline{k}, \underline{k}+1\ldots, K\}$ and her level of influence (on the aggregate level of consumption) $l \in \mathcal{L} = \{\underline{l}, \underline{l}+1, \ldots, L\}$.[4] The joint distribution of susceptibility and influence, denoted by $f \in \Delta(\mathcal{K} \times \mathcal{L})$, is a primitive of the model. Depending on the application we have in mind we may allow for negative influence, but assume that $\mathbb{E}[l] > 0$ to ensure that there are complementarities in consumption on average. Every consumer has a non-negative level of susceptibility, $\underline{k} \geq 0$.

[4]For simplicity, we assume that $\mathcal{K}$ and $\mathcal{L}$ are intervals of integers. Our results extend to the more general case.

For a given level of aggregate consumption $\bar{x}$, consumer i of type (k, l) chooses among the consumption menu the quantity-price pair (x_i, p_i) that maximizes

$$u_i(x_i, p_i, \overline{x}) = (1 + \gamma k \bar{x})\, x_i - \frac{1}{2} x_i^2 - p_i \tag{3.1}$$

The aggregate level of consumption $\bar{x}$ is a weighted average of individual consumption

$$\bar{x} = \sum_{k,l} \frac{l f_{k,l}}{\mathbb{E}[l]} x_{k,l} \tag{3.2}$$

where we denote the consumption of type (k, l) by $x_{k,l}$.[5] We impose a participation constraint with an outside option normalized to zero. The good is produced by a monopolist at zero marginal cost. He offers a menu of contracts $(x, p) = \{(p_{k,l}, x_{k,l})\}_{k,l \in \mathcal{K} \times \mathcal{L}}$ subject to sorting and participation constraints associated to a particular observability assumption. Individual consumption decisions (which are made taking $\overline{x}$ as given) determine the aggregate externality $\overline{x}$. Note that every set of contracts induces a game among the consumers at the consumption stage, as aggregate consumption is endogenous. Still, by the following lemma, restricting the menu to one contract per type is without loss of generality.

Lemma 3.1. *For any observability assumption, there exists a solution to the principal's problem if*

$$\gamma \frac{KL + \mathbb{E}[kl]}{\mathbb{E}[l]} < 1.$$

Furthermore, in the revenue maximal Bayesian equilibrium under the revenue maximal set of contracts, no consumer type randomizes between contracts. The same result holds for welfare maximization.

The intuition behind this result is simple. Suppose a type randomizes between contracts. As the utility function is concave in consumption, she prefers a contract that gives the expected level of consumption and transfer to the principal with certainty. Replacing the contracts she randomizes over with such a contract leaves the expected level of aggregate consumption $\bar{x}$ unchanged. Furthermore, we can increase the transfer to the principal, as to keep utility of the agent the same after the change in the menu of contracts. Hence, she has no changed incentives to deviate to other contracts. Furthermore, the range of allocations that are chosen by consumers of

[5]Equivalently, this is an implicit definition of the level of influence of an agent.

	susceptibility observable	susceptibility not observable
influence observable	$\varnothing$	$\bigcup_{l\in\mathcal{L}} (\mathcal{K} \times l)$
influence not observable	$\bigcup_{k\in\mathcal{K}} (k \times \mathcal{L})$	$(\mathcal{K} \times \mathcal{L})^2$

Table 3.2: The four different observability assumptions as sets of feasible deviations.

susceptibility k shrinks compared to the initial set of contracts. This implies that imitating types of susceptibility k has only gotten less attractive. The sufficiency of sorting constraints at the borders of the range of x chosen by consumers of a given susceptibility is established in the text in the simpler setting without mixing in 3.4.

We will consider several monopolist problems, each correspoding to a different assumption on which consumer charactersistics are observable. Throughout these problems, the objective, the aggregate network effect and the participation constraints will remain the same. Depending on the misreports that are feasible, the problem will have different sorting constraints. Denote the set of feasible deviations by $A \subset (\mathcal{K} \times \mathcal{L})^2$, where $(k,l),(k',l') \in A$ means that type k,l can immitate type k',l' and consequently a feasible allocation must satisfy the sorting constraint

$$(1+\gamma\overline{x}k)\,x_{k,l} - \frac{1}{2}x_{k,l}^2 - p_{k,l} \geq (1+\gamma\overline{x}k)\,x_{k',l'} - \frac{1}{2}x_{k',l'}^2 - p_{k',l'}.$$

With a slight abuse of notation we identify the sorting constraint with the associated (pair of) types. The problem correspoding to a set of feasible deviations A is

$$\prod(A) := \max_{\{p_{k,l},x_{k,l}\}_{k,l\in\mathcal{K}\times\mathcal{L}}} \sum f_{k,l}p_{k,l} \tag{3.3}$$

$$\text{s.t.} \qquad \bar{x} = \sum_{k,l} \frac{l f(l,k)}{\mathbb{E}_f[l]} x_{k,l} \tag{ANE}$$

$$\forall k,l: \quad (1+\gamma\overline{x}k)\,x_{k,l} - \frac{1}{2}x_{k,l}^2 - p_{k,l} \geq 0 \tag{P$_{k,l}$}$$

$$\forall (k,l),(k',l') \in A: \quad (1+\gamma\overline{x}k)\,x_{k,l} - \frac{1}{2}x_{k,l}^2 - p_{k,l} \geq (1+\gamma\overline{x}k)\,x_{k',l'} - \frac{1}{2}x_{k',l'}^2 - p_{k',l'} \tag{IC$_{k,l\to k',l'}$}$$

To save on notation, we suppress the non-negativity constraints $x_{k,l} \geq 0$. table 3.2 specifies the set of feasible deviations associated to each observability assumption

Throughout the paper, we will let ζ denote the Lagrange multiplier associated to the ANE constraint, i.e. marginal revenues associated to an exogenous increase in aggregate consumption.

We allow for a principal's objective function of the form

$$\max_{\overline{x},\{p_{k,l},x_{k,l}\}_{k,l\in\mathcal{K}\times\mathcal{L}}} R(\overline{x}) + \kappa c(\overline{x}) \tag{3.4}$$

where $R(\overline{x})$ is the maximal revenue $\sum f_{k,l}p_{k,l}$ from a screening contract $\{p_{k,l}, x_{k,l}\}_{k,l\in\mathcal{K}\times\mathcal{L}}$ that induces network effect is $\overline{x}$, while c caputures additional social gains or losses associated to aggregate production. Clearly, c would be zero in the application to the sale of the network good, but the generalization to a non-trivial function c allows a more appropriate embedding of the other applications we discuss in the remainder of this section. We assume that c is increasing and concave with the normalization $c''(0) = 1$. In this paper we focus on the shape of the optimal contract (x, p) (under different observability assumptions) and on finding the optimal level of $\bar{x}$ for $\kappa = 0$. Having tackled this problem, the characterization of the solution to the single variable maximization problem (3.4) is a straightforward extension that is covered by the proofs in the Appendix. The structure of the optimal contracts and rents remains unchanged, only the optimal level of $\bar{x}$ is affected.

We now proceed discussing in more detail the application to the sale of network good, as well as alternative economic problems that map into our framework.

3.2.2 Sale of a Network Good

Our main motivation, which we discuss also in the introduction is the sale of a network good.

There are two interpretations of the reduced form (3.1)-(3.2): *local network effects with residual network uncertainty* and *global network effects.* Only the former is derived from an explicit network formation model, which we will discuss in the next paragraph. In the latter, agents directly care about the weighted population average of x, e.g. because of a desire to conform or aggregate network effects of consumption, investment or business activity. Agents differ both in their desire to conform or dependence on aggregate network effects and their intensity of creating network effects for others (visibility, social status or provision of essential services).

Though influence doesn't enter the utility function directly in our specification, the analysis generalizes to this case. If individuals savor influence in a way that is separable from their individual consumption, this is equivalent to a renormalization of the outside option. Gains from influence that come from being invested and central in the diffusion of a popular product enter the utility function just like susceptibility and are captured by a suitable redefinition of the type space. The case where influential consumers have a higher marginal utility from individual consumption that is independent of the aggregate complicates the analysis and is discussed in the conclusion.

We use a network formation model formulated and applied in Galeotti and Goyal (2009); Fainmesser and Galeotti (2016a,b). Consumers are connected by a directed network. When there is a link from i to j we say that consumer i is influenced by agent j. Consumption externalities flow across this network, i.e. a consumer's marginal utility of the good increases as others who influence her increase their consumption. Formally, let I_i be the set of consumers who influence i; the ex-post utility of consumer i is given by

$$u_i\left((x_j)_{j\in[0,1]}, p_i\right) = x_i + \gamma x_i \sum_{j\in I_i} x_j - \frac{1}{2}x_i^2 - p_i \tag{3.5}$$

where γ is the intensity of network effects.

In this interpretation, the influence parameter l coincides with the agent's in-degree, while the susceptibility parameter k is his out-degree. When making their consumption choices, consumers don't know the network structure, but only their in- and out-degree. They take expectations over their realized utility conditional on this information alone.[6] So, the utility can be expressed as

$$u_i(x_i, p_i, \overline{x}) = x_i + \gamma k_i \bar{x} x_i - \frac{1}{2}x_i^2 - p_i \tag{3.6}$$

where

$$\mathbb{E}\left[x_j | j \in I_i\right] = \sum_{k,l} \frac{l f_{k,l}}{\mathbb{E}[l]} x_{k,l} = \bar{x} \tag{3.7}$$

independent of i. When forming expectations, individuals take account of the fact that they are more likely to link to influential individuals which consequently need to be over-counted relative to their frequency in determining the expected consumption of a neighbor.[7] Clearly, equations (3.6) and (3.7) coincide with (3.1)-(3.2).

[6]Formally, we model the network formation as follows: There is a unit interval of consumers, ordered by in-degree l. Denote the in-degree of consumers at $i \in [0,1]$ by $l(i)$, an increasing step function with finite range. After consumption decisions are made, a consumer with out-degree k draws k consumers independently from the unit interval with density $\frac{l(i)}{\mathbb{E}[l]}$ and links to them. Every consumer is drawn and linked to by l other consumers.

[7]For discussion of further effects of this "friends paradox", see Jackson (2017).

3.2.3 Big-push development (Made-in model)

Consider the problem of a ministry of commerce that wants to foster economic development in an area. There are spillovers between firms. These spillover come from the external effects in the provision of investment goods and human capital (economic activity at other firms provides a thick market for such goods and skilled workers which lowers the cost of production). Different firms are heterogeneous in how their activity impacts the aggregate shifter (the influence parameter l, e.g. machine tools firms provide higher investment externalities than textile industry) and on how dependent they are on the externality (the susceptibility parameter k). The regulator then wants to incentivize production from influential sectors to increase total revenues, and we allow for additional effects through the c function.

Equivalently, consider a model of demand spill-overs. Firms produce for export. The aggregate shifter represents the degree to which the region has an image for producing high quality or fashionable products (made-in effect). Different firms are heterogeneous in how their activity impacts this aggregate shifter (luxory sectors like fashion and design are more effective at establishing a country's reputation) and on how dependent they are this perception (products that are clearly recognizable like clothing are more susceptible than commodities). The regulator then wants to incentivize high-quality production from influential sectors to increase total revenues.

3.2.4 Nonlinear taxation of externality producing goods

Firms produce goods using a polluting process. Firms differ both in how their level of production affects total pollution (for instance, they run a more or less polluting plant), and in how much aggregate pollution affects their profit (for instance, they are differently exposed to the consequence of climate change, both in terms of their marginal cost of production and demand). A regulator controlling the aggregate level of externality while rasing tax revenues needs to take into account these two dimensions of heterogeneity.

Our framework offers a natural environment to study this problem. The influence l of a firm measures the amount of pollution per unit produced, while k measures the impact of aggregate pollution on profits per unit produced. In this case of negative externalities, we set $\gamma < 0$ and $\kappa < 0$. Our results stated in the text for $\gamma > 0$ generalize naturally.[8] The planner chooses a tax schedule to maximize a weighted

[8] With $\gamma < 0$, the order of lexicographic monotonicity is inverted. Output is decreasing in influence and susceptibility.

sum of revenues (taking into account the effect of pollution on profitability) and the costs of pollution to society $c(\bar{x})$.

3.2.5 Price setting among gross substitutes firms

Firms selling gross substitutes simultaneously choose their price subject to the linear demand schedule

$$D\left(p,\overline{p}\right)=\left(1+\gamma\overline{p}k-\frac{1}{2}p\right)$$

where the aggregate price level is $\overline{p}=\sum_{k,l}\frac{lf_{k,l}}{\mathbb{E}[l]}p_{k,l}$. Firm's influence l parametrizes the weight of its price in the aggregate price level, while k measures its degree of substitutability relative to the price aggregator. A firm choosing p to maximize $D\left(p,\overline{p}\right)\cdot p$ clearly induces the utility specification (3.1), but we lack a microfoundation for the aggregate price index.

3.3 Efficient Allocation

We begin establishng our efficiency benchmark which, as is standard, coincides with the allocation that realizes under perfect information. So the efficient allocation is characterized by the solution of $\Pi\left(\varnothing\right)$, which next subsection characterizes.

3.3.1 The First Best

The following Proposition characterizes full-information allocations, *unit* prices, aggregate statistics and revenues.

Proposition 3.1. *The optimal menu of contracts* $\{(p_{k,l},x_{k,l})\}_{k,l\in\mathcal{K}\times\mathcal{L}}$ *under full information is such that*

$$x_{kl}=1+\gamma\bar{x}k+\frac{l}{\mathbb{E}[l]}\zeta \tag{3.8}$$

$$\frac{p_{k,l}}{x_{k,l}}=\frac{1}{2}\left(1+\gamma\overline{x}k-\frac{l}{\mathbb{E}[l]}\zeta\right) \tag{3.9}$$

where

$$\bar{x}=\frac{1+\frac{\mathbb{E}[l^2]}{\mathbb{E}[l]^2\left[1-\gamma\frac{\mathbb{E}[kl]}{\mathbb{E}[l]}\right]}\gamma\mathbb{E}[k]}{1-\gamma\frac{\mathbb{E}[kl]}{\mathbb{E}[l]}-\frac{\gamma^2\mathbb{E}[l^2]\mathbb{E}[k^2]}{\mathbb{E}[l]^2\left[1-\gamma\frac{\mathbb{E}[kl]}{\mathbb{E}[l]}\right]}},\qquad \zeta=\sum_{k,l}f_{kl}\gamma kx_{kl}=\frac{\gamma\mathbb{E}[k]+\frac{\gamma^2\mathbb{E}[k^2]}{\left[1-\gamma\frac{\mathbb{E}[kl]}{\mathbb{E}[l]}\right]}}{1-\gamma\frac{\mathbb{E}[kl]}{\mathbb{E}[l]}-\frac{\gamma^2\mathbb{E}[l^2]\mathbb{E}[k^2]}{\mathbb{E}[l]^2\left[1-\gamma\frac{\mathbb{E}[kl]}{\mathbb{E}[l]}\right]}}$$

The principal extracts all surplus as revenue

$$R = \frac{1}{2} + \gamma\bar{x}\mathbb{E}[k] + \frac{1}{2}\gamma^2\bar{x}^2\mathbb{E}[k^2] - \frac{1}{2}\zeta^2\frac{\mathbb{E}[l^2]}{\mathbb{E}[l]^2}. \tag{3.10}$$

The efficient consumption is the optimal consumption for given externality $\bar{x}$, $1+\gamma\bar{x}k$, with the adjustment term $\frac{l}{\mathbb{E}[l]}\zeta$ taking into account the spillover through aggregate consumption $\bar{x}$. This adjustment is proportional to the marginal value of $\bar{x}$, ζ which is proved to be positive, and to the ratio between the weight of type (k, l) in social welfare f_{kl} and its weight in determining $\bar{x}$, $\frac{lf(l,k)}{\mathbb{E}_f[l]}$.

As the monopolist extracts all surplus, agents receive the same level of utility (zero) in the optimal contract. In particular, influence is not rewarded. Highly influential agents over-consume (which they dislike due to the quadratic term), though they are compensated with a quantity discount in the form of lower unit price.

3.3.2 Observable Susceptibility

We now turn to the case where to monopolist observes how susceptible a consumer is, but a consumer's influence is unobservable. The revenue maximization problem needs to satisfy the sorting constraints $\bigcup_{k\in\mathcal{K}}(k\times\mathcal{L})$,[9] namely for the pairs of types that share the same k.

First, note that l does not directly enter the utility function. Consequently, the relevant set of incentive constraint is equivalent to requiring that all consumers with a given k have the same utility in their assigned contract; should this condition fail, every consumer with a given k would mimick the type k, l' whose contract delivers the highest level of utility. Formally,

Lemma 3.2. *A menu of contracts satisfies the $\bigcup_{k\in\mathcal{K}}(k\times\mathcal{L})$ sorting constraints if and only if the utility of type (k, l) is independent of the level of influence l.*

[9]Recall we identified the set of revelant sorting constraint with the indeces pairs that are involved.

Proof. Fix an arbitrary k and suppose the set of contracts $\{x_{k,l}, p_{k,l}\}_{l\in\mathcal{L}}$ delivers the same utility $u_{k,l} = (1+\gamma\overline{x}k)\,x_{k,l} - \frac{1}{2}x_{k,l}^2 - p_{k,l}$ for all $l \in \mathcal{L}$. Clearly, there is no incentive to misrepresent your type.

That this is necessary is immediate from $\text{IC}_{k,l\to k,l'}$and $\text{IC}_{k,l'\to k,l}$:

$$\underbrace{(1+\gamma\overline{x}k)\,x_{k,l} - \frac{1}{2}x_{k,l}^2 - p_{k,l}}_{u_{k,l}} \geq \underbrace{(1+\gamma\overline{x}k)\,x_{k,l'} - \frac{1}{2}x_{k,l'}^2 - p_{k,l'}}_{u_{k,l'}} \geq \underbrace{(1+\gamma\overline{x}k)\,x_{k,l} - \frac{1}{2}x_{k,l}^2 - p_{k,l}}_{u_{k,l}}$$

□

The intuition is that influence doesn't interact with the contract terms, so it can not introduce distortions in the form of information rents: Even though the problem has a full dimension of incomplete information, it collapses for given k. The first best contract satisifes the $l-$independece condition since all types receive zero utility, so by lemma 3.2 above it satisfies the incentive constraints that are relevant for the observable suscepitbility case. Its optimality is then immediate.

Proposition 3.2. *The problem with known susceptibility k is equivalent to the problem under full information. The optimal menu of contracts is as described in Proposition 3.1.*

By 3.2 observable susceptibility induces the same allocation as the full information first best benchmark, that features no information rents. This happens because the principal observes the only dimension in which she can actively screen. Eliciting influence does not create any rents by itself.

3.4 Private Information

We now turn to the more interesting full screening problem, i.e. the case in in which both influence and susceptibility are not observed by the principal. She then solves the 2-dimensional screening problem with consumption externality $\prod\left((\mathcal{K}\times\mathcal{L})^2\right)$. The aim of this section is to simplify and solve this problem, which also constitutes the main technical contribution of this paper. The following proposition characterizes properties of the optimal allocation

Proposition 3.3. *A menu of contracts* $\{(x_{kl}, p_{k,l})\}_{k,l\in\mathcal{K}\times\mathcal{L}}$ *is optimal in* $\prod\left((\mathcal{K}\times\mathcal{L})^2\right)$ *only if*

1. The utility of type (k,l) is independent of the level of influence l. That is, for each k, l, l',

$$(1+\gamma\overline{x}k)\, x_{k,l} - \frac{1}{2}x_{k,l}^2 - p_{k,l} = (1+\gamma\overline{x}k)\, x_{k,l'} - \frac{1}{2}x_{k,l'}^2 - p_{k,l'} \tag{H}$$

   (a) It satisfies k-monotonicity, that is for every k, k', l, l',

$$(x_{k,l} - x_{k',l'})\,(k - k') \geq 0 \tag{M}$$

   (b) The participation constraint is binding for only the lowest susceptibility types. That is,

$$(1+\gamma\overline{x}\underline{k})\, x_{\underline{k},l} - \frac{1}{2}x_{\underline{k},l}^2 - p_{\underline{k},l} = 0. \tag{P}$$

We allocations satisfying H, P, and M **admissible**.

Since the problem contains all incentive compatibility constraints along the influence dimension $\left(\bigcup_{k\in\mathcal{K}}(k\times\mathcal{L})\subset(\mathcal{K}\times\mathcal{L})^2\right)$, Lemma lemma 3.2 implies that utility has to be constant along the l dimension, which is point 1. Point 2 is a standard monotonicity result derived by comparing incentive constraints: types with high marginal utility need to get a high level of consumption. Finally, by the sorting constraints, it is sufficient to consider the participation constraint of type $\underline{k}$ as all other participation constraints are implied.

We further reduce the cardinality of the sorting constraints by noticing that some of them are redundant. Given $l-$invariance of utility, a slice $k\times\mathcal{L}$ of the type space can be treated as a single type for the purpose of outward deviations. In addition, we can rank the attractivenes of contracts in each $k\times\mathcal{L}$ slice by the level of the allocation: Higher susceptibility types will prefer the highest allocation contract, while lower susceptibility types will prefer the lowest allocation contract. Consequently, for types in $k \times \mathcal{L}$, the relevant downward deviation is towards the contract giving the highest consumption in the $k-1 \times \mathcal{L}$ slice, whereas the relevant upward deviation is towards the contract giving the lowest consumption in the $k+1 \times \mathcal{L}$ slice.

Definition 3.1. Fix a menu of contracts $\{(x_{kl}, p_{k,l})\}_{kl \in \mathcal{K}\times\mathcal{L}}$ and, for each k, pick

$$l_k \in \arg\min_{\widetilde{l}} x_{k,\widetilde{l}}, l^k \in \arg\max_{\widetilde{l}} x_{k,\widetilde{l}}$$

We call the set of constraints

$$\left\{\left((k,l),(k-1,l^{k-1})\right)\right\}_{k>\underline{k},l\in\mathcal{L}} \cup \{((k,l),(k+1,l_{k+1}))\}_{k<K,l\in\mathcal{L}} \subset (\mathcal{K}\times\mathcal{L})^2$$

a set of **extremal sorting (ES) constraints**.

To clarify the definition, Figure 3.1 below depicts the set of ES constraints for an arbitrary proposed consumption allocation.

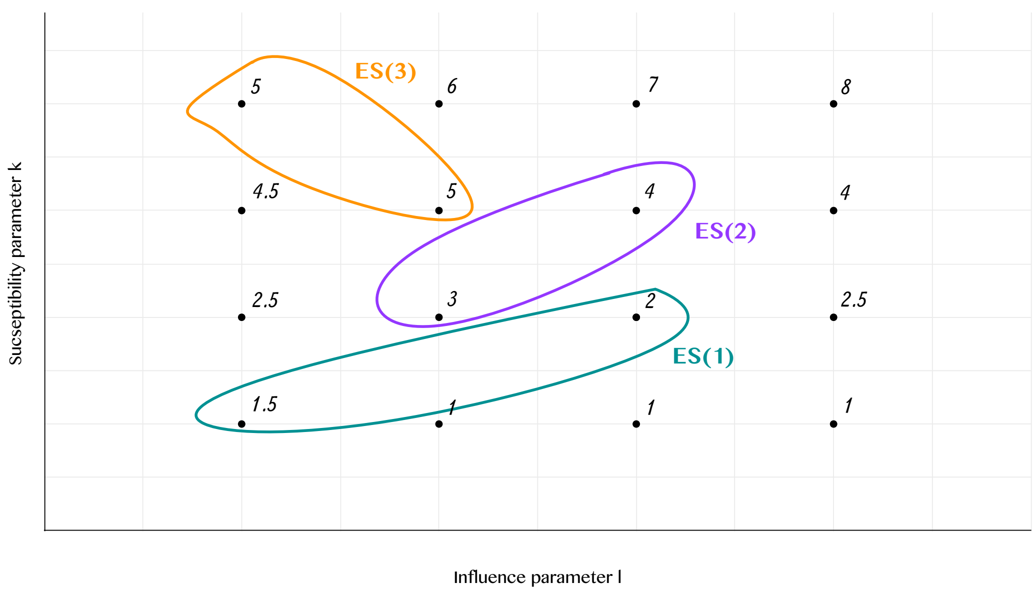

Figure 3.1: Extremal Sorting constraints for a proposed consumption allocation

If it is not profitable to deviate to the contract with the largest (smallest) level of consumption in the slice, it isn't profitable to deviation into the slice at all. This leads to the following

Proposition 3.4. *An admissible allocation satisfying a set of ES constraints satisfies all sorting constraints.*

Proof. Downward: consider the sorting constraint from type k, l to type k', l', where $k > k'$. Then

$$\begin{aligned} u_{kl} = u_{kl_k} &\geq u_{k-1l^{k-1}} + \gamma\bar{x}x_{k-1,l^{k-1}} \\ &\geq u_{k-1l_{k-1}} + \gamma\bar{x}x_{k-1,l^{k-1}} \end{aligned}$$

where the first inequality is the local extremal downward sorting constraint and the equalities follow because the utility under compliance with the menu is independent of l. Repeating the above argument, we arrive at

$$\begin{aligned} u_{kl} &\geq u_{k'l^{k'}} + \gamma\bar{x}\sum_{i=k'}^{k-1} x_{il^i} \\ &= u_{k'l'} + \gamma\bar{x}\sum_{i=k'}^{k-1} x_{il^i} \\ &\geq u_{k'l'} + \gamma\bar{x}\sum_{i=k'}^{k-1} x_{k'l'} \\ &= \left(1 + \gamma\bar{x}k\right) x_{k',l'} - \frac{1}{2}x_{k',l'}^2 - R_{k',l'} \end{aligned}$$

where the last inequality follows from k-monotonicity (Lemma 3.3) together with the definition of l^i as the level of influence maximizing x_{il^i}. Note that the last term is just the utility of type k, l pretending to be type k', l', which is not profitable, as we wanted to show.

An analogous argument works for upwards deviations.

□

The next step is to identify the extremal types within a slice $k \times \mathcal{L}$. We show that x_{kl} is not only k-monotonic, but increasing with respect to the lexicographic order $\succ_L$ on $\mathcal{K} \times \mathcal{L}$ where $\mathcal{K}$ is the dominant dimension

$$(k, l) \succ_L (k', l') \iff k > k' \text{ or } k = k', l > l'$$

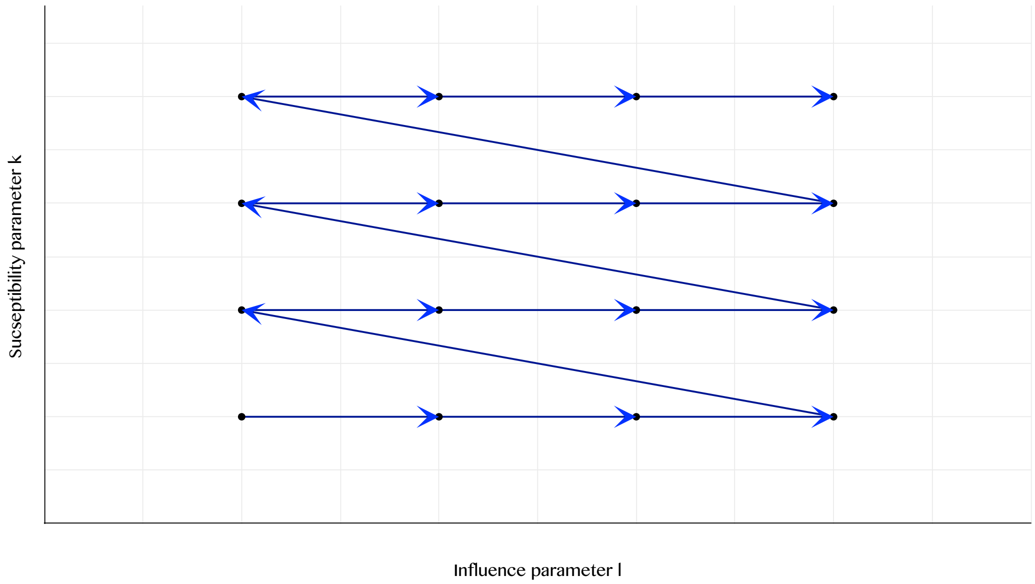

Figure 3.2: Lexicographic monotonicity

Theorem 3.1 (Lexicographic Monotonicity). *In the solution to GP, x_{kl} is weakly increasing in the $\succ_L$ order on $\mathcal{K} \times \mathcal{L}$.*

The argument goes as follows: for each k, the unconstrained allocation is monotonic in l. Since 3.4 implies that only the largest $x_{k,l}$ for a slice $k \times \mathcal{L}$ is relevant for deviation, adding the sorting constraints does not alter this property. Combining this fact with k- monotonicity implies that the optimal allocation has to be lexicographic monotonic – either increasing in the second component or decreasing in the second component. The direction of the order depends on the sign of the multiplier ζ in the optimal allocation. We can show that the marginal value of aggregate consumption ζ has to be positive, which selects the $\succ_L$ lexicographic order on $\mathcal{K} \times \mathcal{L}$.[10]

A positive marginal value of aggregate consumption for the monopolist is a natural though not immediate result. In contrast to the symmetric information benchmark, under asymmetric information aggregate consumption $\bar{x}$ impacts revenues in two opposing ways: On the one hand, increasing $\bar{x}$ increases total surplus; on the other hand, it increases the information rents paid to consumers. We show that the first force dominates. Hence consumption is increasing in l: this increases aggregate consumption $\bar{x}$, counteracting some of the downward distortion due to screening.

[10]Note that by contrast to $k-$monotonicity, lexicographic monotonicity doesn't follow from the constraints alone, but is based on optimality considerations.

3.4.1 The Relaxed Problem

Using the results derived in the previous section, we can rewrite the principal's problem as a monotonicity constrained optimization in terms of virtual values. This problem is one-dimensional along the lexicographic order.

Proposition 3.5. *The problem $\prod\left((\mathcal{K}\times\mathcal{L})^2\right)$ is equivalent to*

$$\max_{\{x_{k,l}\}_{k,l\in\mathcal{K}\times\mathcal{L}}} \sum f_{k,l}\left[(1+\gamma\bar{x}k)\,x_{k,l}-\frac{1}{2}x_{k,l}^2-\chi_{l=L}\left\{\frac{1-F_k}{f_{kl}}x_{k,l}\right\}\right] \tag{UP}$$

$$s.t.\qquad \bar{x}=\sum_{k,l}\frac{lf(l,k)}{\mathbb{E}_f[l]}x_{k,l} \tag{ζ}$$

$$\forall k,l\in\mathcal{K}\times\{\underline{l}+1,\dots,L\}:\quad x_{k,l}\geq x_{k,l-1}\ l>\underline{l} \tag{3.11}$$

$$\forall k\in\{\underline{k}+1,\dots,K\}:\quad x_{k,\underline{l}}\geq x_{k-1,L} \tag{3.12}$$

$$x_{\underline{k},\underline{l}\geq}0 \tag{3.13}$$

The remaining nonstandard feature is the externality through the endogenous aggregate quantities $(\bar{x},\zeta)$. We will solve the problem for fixed $(\bar{x},\zeta)$ and then solve for the aggregate quantities.

To solve the problem for given $(\bar{x},\zeta)$, we rewrite the objective function of (UP) as

$$\sum f_{k,l}\left[\left(\underbrace{1+\gamma\bar{x}\left(k-\chi_{l=L}\left\{\frac{1-F_k}{f_{kl}}\right\}\right)+\zeta\frac{l}{\mathbb{E}[l]}}_{:=g(k,l)}\right)x_{k,l}-\frac{1}{2}x_{k,l}^2\right]$$

and apply ironing. In this problem, ironing is typically necessary and cannot be avoided by assuming a natural condition as e.g. a monotone hazard rate. By the nature of our problem, there are two sources of monotonicity violations in the virtual value. First, sorting constraints only affect types with $l=L$ directly. The resulting downward distortion will typically be propagated by the monotonicity constraints. This happens whenever

$$\zeta\frac{1}{\mathbb{E}[l]}<\gamma\bar{x}\frac{1-F_k}{f_{k,L}},$$

i.e. whenever the downward distortion due to sorting constraints is larger than the difference in efficient consumption due to lower influence. Unless the highest influence type is very likely ($f_{k,L}$ large), this condition will be met. The second source of violations of monotonic virtual values is the jump between type k, L and $k+1, \underline{l}$: One the one hand, the latter has a higher valuation because he is more susceptible to influence, on the other hand, he is less influential, which depresses consumption. There will be a non-monotonicity if

$$\gamma\bar{x}\left(1+\frac{1-F_k}{f_{k,L}}\right) < \zeta\frac{L-\underline{l}}{\mathbb{E}[l]}$$

Combining both inequalities, we see that we require ironing unless for all k

$$\zeta\frac{L-\underline{l}}{\mathbb{E}[l]} < \gamma\bar{x}\left(1+\frac{1-F_k}{f_{k,L}}\right) < \gamma\bar{x}+\zeta\frac{1}{\mathbb{E}[l]}$$

which involves bounds on endogenous quantities.

Now, we turn to the ironing procedure. The problem is one-dimensional in the lexicographic order and hence we can apply standard techniques. We will refer to this order by a subscript "lex" to avoid confusion. Let $(k^*, l^*)(q) := \min_{lex}\left\{(k,l) \mid \sum_{k',l' \leq_{lex} k,l} f_{k'l'} \geq q\right\}$ denote the type at the $q-$quantile for any $q \in [0,1]$. The cumulative virtual value is given by

$$G(q) = \sum_{(k',l')<_{lex}(k^*,l^*)(q)} f_{k',l'} g(k',l') + \left(q - \sum_{k',l' \leq_{lex}(k^*,l^*)(q)} f_{k'l'}\right) g((k^*,l^*)(q)).$$

It follows from Myerson (1981); Toikka (2011) that ironing the original problem is equivalent to convexifying G. Let $H := \text{Conv}(G) := \max\{h \leq G | h \text{ is convex}\}$. If $H(q) < G(q)$, the monotonicity constraints are active at the corresponding type and there is bunching.

Lemma 3.3. *H induces a partition of types, B. The partition cells are ordered by the lexicographic order and the optimal allocation for a given $(\bar{x}, \zeta)$ is constant within cells and strictly increasing across cells.*

We can hence write x_b for the quantity in partition cell b and arrive at the allocation as a linear function of the aggregate control and multiplier $(\bar{x}, \zeta)$. For

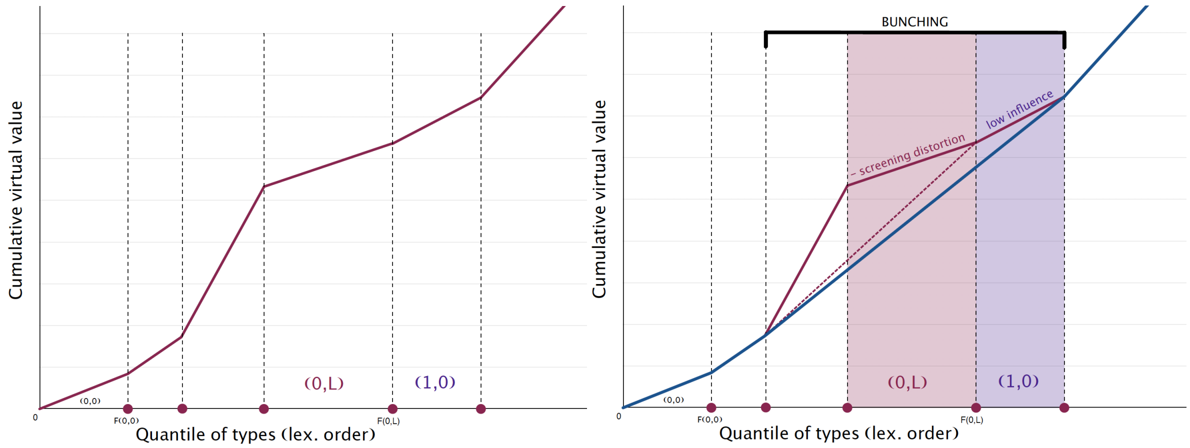

Figure 3.3: Typical ironing region

$x_{k,l}$ with $(k,l) \in b \in B$, we get

$$x^*_{k,l} = x_b = \max\left\{1 + \gamma\overline{x}\left(\mathbb{E}\left[\boldsymbol{k}\,|b\right] - \frac{\sum_{\left\{k:\left(k,\bar{l}\right)\in b\right\}}\left(1-F_k\right)}{\sum_{(k,l)\in b} f_{k,l}}\right) + \zeta\frac{\mathbb{E}\left[\boldsymbol{l}\,|b\right]}{\mathbb{E}\left[\boldsymbol{l}\right]}, 0\right\},$$

where the max ensures that $x_b \geq 0$.

Remark 3.1. As our type space is finite, the convexification is easy to compute. A simple algorithm proceeds downwards in the lexicographic order and "greedily" irons out violations of convexity as it encounters them. It finishes in at most $|\mathcal{K} \times \mathcal{L}| + 1$ steps.

Define a self-map $\Gamma : \mathbb{R}^2_+ \mapsto \mathbb{R}^2_+$ by

$$\Gamma(\bar{x},\zeta) = \left(\sum \frac{l f_{k,l}}{\mathbb{E}[l]} x^*_{k,l}(\bar{x},\zeta), \gamma\left\{\sum_{k,l} f_{k,l} k x^*_{k,l}(\bar{x},\zeta) - \sum_{k=\underline{k}}^{K-1}\left(1-F_k\right) x^*_{k,L}(\bar{x},\zeta)\right\}\right)$$

This mapping updates the level of aggregate consumption and its multiplier based the optimal allocation at an initial guess.

Lemma 3.4. Γ *is continous and has a precompact range if* $1 > \gamma\frac{L}{\mathbb{E}[l]}\left(K + \mathbb{E}[k]\right)$.[11]

[11] This is one way to get an upper bound. We conjecture, that it is not tight and that $1 > \gamma\frac{KL}{\mathbb{E}[l]}$ is sufficient.

Hence, there exists a fixed point.

Theorem 3.2. *If* $1 > \gamma\frac{L}{\mathbb{E}[l]}\left(K + \mathbb{E}[k]\right)$*, then there exists a solution to*

$$\Gamma(\bar{x}^\star, \zeta^\star) = (\bar{x}^\star, \zeta^\star). \tag{3.14}$$

Furthermore, the optimal menu of contracts solving $\prod\left((\mathcal{K}\times\mathcal{L})^2\right)$ *is* $\left\{\left(x^*_{k,l}(\bar{x}^\star,\zeta^\star), p^*_{k,l}(\bar{x}^\star,\zeta^\star)\right)\right\}_{k,l\in\mathcal{K}\times}$ *where* $(\bar{x}^\star,\zeta^\star)$ *solves (3.14).*

3.5 Observable Influence

In this section, we analyse the problem when an agent's influence is observable, but his level of susceptibility to the network effect is private information.

3.5.1 Observable Influence, Unobservable Susceptibility

The problem of the monopolist is

$$\Pi\left(\bigcup_{l\in\mathcal{L}}(\mathcal{K}\times l)\right) = \max_{\{p_{k,l},x_{k,l}\}_{k,l\in\mathcal{K}\times\mathcal{L}}} \sum_{k,l} f_{k,l}p_{k,l} \tag{LP}$$

$$\text{s.t.} \qquad \bar{x} = \sum_{k,l}\frac{lf(l,k)}{\mathbb{E}_f[l]}x_{k,l} \tag{3.15}$$

$$\forall l,k,k': \quad (1+\gamma\bar{x}k)\,x_{k,l} - \frac{1}{2}x_{k,l}^2 - p_{k,l} \geq (1+\gamma\bar{x}k)\,x_{k',l} - \frac{1}{2}x_{k',l}^2 - p_{k',l} \tag{IC-$K_{l,k,k'}$}$$

$$\forall k,l: \quad (1+\gamma\bar{x}k)\,x_{k,l} - \frac{1}{2}x_{k,l}^2 - p_{k,l} \geq 0 \tag{P-$L_{k,l}$}$$

Note that the sorting constraints now apply for fixed l, so the problem is one-dimensional. Hence, following the usual arguments, the incentive constraints imply monotonicity and we can rewrite the problem as the maximization of virtual value subject to this monotonicity constraint. This transformation applies for every fixed l, but the optimization problem doesn't factor, as the components for different l are coupled through aggregate consumption $\bar{x}$. Let $F_k^l = \sum_{i=\underline{k}}^k f_{i,l}$ denote the conditional cdf.

Proposition 3.6. *The maximization problem LP is equivalent to*

$$\max_{\{p_{k,l},x_{k,l}\}_{k,l\in\mathcal{K}\times\mathcal{L}}} \sum_{k,l} f_{k,l}\left\{\left(1+\gamma\bar{x}\left(k-\frac{F_K^l-F_k^l}{f_{kl}}\right)\right)x_{kl}-\frac{1}{2}x_{kl}^2\right\} \tag{LPs}$$

$$s.t. \qquad \bar{x}=\sum_{k,l}\frac{l f_{k,l}}{\mathbb{E}[l]}x_{k,l} \tag{3.16}$$

$$\forall k>k', l: \quad x_{k,l}\geq x_{k',l}\geq 0 \tag{3.17}$$

We will assume that virtual values are increasing and hence the monotonicity constraints will be slack.

Assumption 3.1. *The virtual value* $k-\frac{F_K^l-F_k^l}{f_{kl}}$ *is increasing in* k*.*

Analogous to the full information case, the first-order conditions of this problem have two components. The first part is the familiar screening formula, the second adjusts consumption upward for influential individuals in order to provide a stronger network effect.

$$x_{kl}=\max\left\{\underbrace{1+\gamma\bar{x}\left(k-\frac{F_K^l-F_k^l}{f_{kl}}\right)}_{\text{optimal screening for fixed }\bar{x}}+\underbrace{\frac{l}{\mathbb{E}[l]}\zeta}_{\text{provide public good }\bar{x}},0\right\} \tag{3.18}$$

$$\zeta=\gamma\sum_{k,l}f_{kl}\left\{k-\frac{F_K^l-F_k^l}{f_{kl}}\right\}x_{kl} \tag{3.19}$$

$$\bar{x}=\sum_{k,l}f_{kl}l x_{kl} \tag{3.20}$$

Agents receive information rents for their level of susceptibility k. The extent of these rents depends on the level of consumption optimally given to types with the same l and lower k. Therefore, the rent of type k,l is dependent on his (observable) level of influence. We say that there are rents from influence if, for fixed k, the information rent is increasing in l.

Proposition 3.7. *Influential individuals gain an rent if* $\frac{l}{\mathbb{E}[l]}\gamma\Xi - \frac{F_K^l - F_k^l}{f_{kl}}$ *is increasing in* l*, where*

$$\Xi = \frac{\zeta}{\gamma^2\bar{x}} = \sum_l \sum_{\underline{k}}^{K} \left(F_K^l - F_k^l\right)(k+1)\left[1 - \frac{F_K^l - F_{k+1}^l}{f_{k+1,l}} + \frac{F_K^l - F_k^l}{f_{kl}}\right] > 0$$

Influence affects information rents through two channels. First, more influential individuals consume more and high levels of consumption cause high rents. Second, influence has an effect on the downward distortion of consumption by the monopolist. If susceptibility and influence are affiliated, high influence makes it more likely that the agent also has high susceptibility and the monopolist distorts consumption downwards more. The outcome depends on the balance of these two forces. If k and l are independent, only the first is active and influential agents get higher information rents. In general, the condition puts an upper bound on the degree of theiraffilliation. –(But note that this is only a rough intuition, as affiliation also enters Ξ.)–

Comparing the results in this section to the previous one, we see that influence affects an agents utility *only if it is observable and only indirectly, through its impact on susceptibility rents.* In the case of positive consumption externalities, influencers don't get rewarded through nonlinear pricing (though casual observation suggests that there are rewards through other channels). In the case of polution, higly polluting industries are not punished indirectly through a revenue maximizing tax.

3.6 Restrictions on Misreports

It is sometimes natural to assume that agents cannot underreport (or overreport) their susceptibility to or influence on externalities. Take for example consumption externalities mediated through a social network. If the network provider takes effective measures against the purchase of fake subscribers, influencers cannot overreport their number of followers. They could, however, underreport their influence by splitting their activities across different networks and channels. Conversely, if susceptibility corresponds to the time spent on the social network, users cannot underreport it, but they could easily generate fake traffic and thereby overreport their activity. In the case of pollution, it may be possible to provide verifiable evidence on its impact on one's production process. If such evidence is required, firms cannot overreport their susceptibility, but they could still underreport it by not disclosing the evidence.

Polluters could easily prove their activity, but it is much more involved to prove the absence of polluting activities.

In this section, we show that such restrictions on misreports map into the cases analyzed above. Agents always desire to underreport their susceptibility. Hence, if agents can only overreport, this is equivalent to observable susceptibility and the monopolist can implement the first-best. If the condition of Assumption 3.1 and Proposition 3.7 are met, agents always want to overreport their influence. Hence, if doing so is impossible, it is equivalent to observable influence and the optimal contracts are as in the previous section.

3.6.1 No Downward Misreports of Susceptibility

Suppose agents cannot underreport their susceptibility to the externality. We don't restrict their ability to exagerate their susceptibility or misreports of their influence. This is a reasonable assumption in social networks if susceptibility is tightly linked to usage of the network which is identifyable by the network provider and cannot be hidden.

Consider the first-best allocation. The participation constraint is binding for every type. Suppose an agent with type k, l deviates to k', l' with $k' > k$. The utility under this deviation is

$$\hat{u} = u_{k',l'} + \gamma\bar{x}\left(k' - k\right)x_{k',l'} = \gamma\bar{x}\left(k' - k\right)x_{k',l'} < 0 = u_{k,l}$$

Hence, this deviation is not profitable. As the utility is flat along the l dimension, there is also no incentive to misrepresent only influence. Hence, the first-best allocation is incentive compatible and we have the following

Proposition 3.8. *Suppose agents with susceptibility k can only report $\hat{k} \geq k$. Then, the first-best allocation is implementable.*

3.6.2 No Upward Deviations of Influence

Suppose instead that agents can freely underreport their level of susceptibility, but cannot overreport their influence. This is a reasonable assumption in many social media platforms, whereit is comparatively easy for agents to appear less well connected than they truly are. Users can split their activities on the platform by using multiple accounts or using private browsing for some of their interactions. By contrast, it can be difficult to appear more connected than you truly are. While followers on many platforms can be bought, these tend to be "bot accounts" which can often be detected by the site's algorithms.

In this case, agents would desire to underreport their susceptibility, so these deviations have to be taken care off by sorting constraints. Consider the allocation solving LPs, where we assume that influence is observable but susceptibility is not. If the condition of Assumption 3.1 and Proposition 3.7 are met, agents always desire to overreport their level of influence. Since this is impossible, the same allocation is implelemtable, even though influence is not fully observed.

Proposition 3.9. *If $k - \frac{F_K^l - F_k^l}{f_{kl}}$ is increasing in k and $\frac{l}{\mathbb{E}[l]}\gamma\Xi - \frac{F_K^l - F_k^l}{f_{kl}}$ is increasing in l, the solution to the monopolist's problem with underreports and the problem with known influence and unobserved susceptibilit are the same.*

The intuition behind this result is simple. Under the given conditions, consumers always want to underreport their susceptibility and overreport their influence. They hence report the highest feasible level of influence, which is their true level, but need to be encouraged to report their true level of susceptibility through information rents.

3.7 Conclusion

We analyze a screening problem with externalities. Agents have private information about their susceptibility to and influence on the externality. A monopolist principal provides a menu of actions to maximize revenue. Several problems fall into this framework, for example a monopolist firm using nonlinear pricing when there are consumption externalities or a government maximizing tax revenues when there are externalities between firms, positive through external economies of scale or negative through pollution. Even though the problem is two-dimensional at the surface and contracts are linked "globally" trhough the externality, we show that it is nevertheless tractable. The principal screens along the susceptibility dimension while tilting the allocation along the influence dimension to correct for the externality. Eliciting influence is for free: As long as susceptibility is observable, the principal can implement the first-best. If both characteritics are unobservable, we show that the problem can be transformed into a one-dimensional problem along the lexicographic order, increasing in susceptibility as the dominant dimension and either increasing or decreasing in influence, depending on wether positive or negative externalities dominate. There are rents for susceptibility, but no rents for influence. If influence is observable, the problem is equivalent to a family of one-dimensional screening problems coupled through the externality. Influence affects utility only if it is observed and even then only indirectly, through its effect on information rents. Highly influential consumers obtain higher rents if the externality is positive and susceptibility and influence are not too affiliated. Restricting the agents to over- or underreports maps into the observability assumptions we analyze exhaustively.

References

Laurence Ball and David Romer. Real Rigidities and the Non-Neutrality of Money. *The Review of Economic Studies*, 57(2):183–203, 1990. ISSN 0034-6527. doi: 10.2307/2297377.

Jess Benhabib and Roger E. A. Farmer. Indeterminacy and Increasing Returns. *Journal of Economic Theory*, 63(1):19–41, June 1994. ISSN 0022-0531. doi: 10.1006/jeth.1994.1031.

Francis Bloch and Nicolas Quérou. Pricing in social networks. *Games and Economic Behavior*, 80:243–261, July 2013. ISSN 08998256. doi: 10.1016/j.geb.2013.03.006.

John Bryant. A Simple Rational Expectations Keynes-Type Model. *The Quarterly Journal of Economics*, 98(3):525–528, 1983. ISSN 0033-5533. doi: 10.2307/1886025.

Ozan Candogan, Kostas Bimpikis, and Asuman Ozdaglar. Optimal Pricing in Networks with Externalities. *Operations Research*, 60(4):883–905, August 2012. ISSN 0030-364X, 1526-5463. doi: 10.1287/opre.1120.1066.

Russell Cooper and Andrew John. Coordinating Coordination Failures in Keynesian Models. *The Quarterly Journal of Economics*, 103(3):441–463, 1988. ISSN 0033-5533. doi: 10.2307/1885539.

Gergely Csorba. Screening contracts in the presence of positive network effects. *International Journal of Industrial Organization*, 26(1):213–226, January 2008. ISSN 01677187. doi: 10.1016/j.ijindorg.2006.11.003.

Itay P. Fainmesser and Andrea Galeotti. Pricing Network Effects: Competition. Technical report, 2016a.

Itay P. Fainmesser and Andrea Galeotti. Pricing Network Effects. *The Review of Economic Studies*, 83(1):165–198, January 2016b. ISSN 0034-6527, 1467-937X. doi: 10.1093/restud/rdv032.

Joseph Farrell and Garth Saloner. Standardization, Compatibility, and Innovation. *The RAND Journal of Economics*, 16(1):70–83, April 1985. ISSN 0741-6261. doi: 10.2307/2555589.

Andrea Galeotti and Sanjeev Goyal. Influencing the influencers: A theory of strategic diffusion. *The RAND Journal of Economics*, 40(3):509–532, September 2009. ISSN 1756-2171. doi: 10.1111/j.1756-2171.2009.00075.x.

Arne Rogde Gramstad. Nonlinear Pricing with Local Network Effects. Technical report, 2016.

Matthew O. Jackson. The Friendship Paradox and Systematic Biases in Perceptions and Social Norms. SSRN Scholarly Paper ID 2780003, Social Science Research Network, Rochester, NY, November 2017.

Ali Jadbabaie and Ali Kakhbod. Optimal Contracting in Networks. *SSRN Electronic Journal*, 2016. ISSN 1556-5068. doi: 10.2139/ssrn.2774600.

Michael L. Katz and Carl Shapiro. Network Externalities, Competition, and Compatibility. *The American Economic Review*, 75(3):424–440, June 1985. ISSN 0002-8282.

Roger B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6 (1):58–73, 1981. ISSN 0364-765X.

Ilya Segal. Contracting with Externalities. *The Quarterly Journal of Economics*, 114 (2):337–388, 1999. ISSN 0033-5533.

Fanqi Shi and Yiqing Xing. Screening with Network Externalities. Technical report, 2018.

Arun Sundararajan. Nonlinear pricing and type-dependent network effects. *Economics Letters*, 83(1):107–113, April 2004. ISSN 0165-1765. doi: 10.1016/j.econlet.2003.10.009.

Juuso Toikka. Ironing without control. *Journal of Economic Theory*, 146(6):2510–2526, November 2011. ISSN 00220531. doi: 10.1016/j.jet.2011.06.003.

Yang Zhang and Ying-Ju Chen. Optimal Nonlinear Pricing in Social Networks under Asymmetric Network Information. SSRN Scholarly Paper ID 2876087, Social Science Research Network, Rochester, NY, October 2017.

Appendix

Proof of Lemma 3.1: Suppose towards a contradiction that there is mixing in the optimum. Consider a type k, l playing a non-degenerate mixed strategy $\sigma_{k,l}$ with finite support. Let $x'_{k,l} = \mathbb{E}_{\sigma_{k,l}} x$ and $R'_{k,l} = \mathbb{E}_{\sigma_{k,l}} R + \frac{1}{2}\mathbb{E}_{\sigma_{k,l}} x^2$. Remove all contracts that were only in the support of $\sigma_{k,l}$ and not played by any other type and add the contract $(x'_{k,l}, R'_{k,l})$ to the menu. We propose that it is an equilibrium for all other types to follow the same strategy as in the original equilibrium and for type k, l to choose the new contract with certainty. Note that with this play, $\bar{x}$ is unchanged from the old candidate equilibrium.

The described strategy combination is feasible as all contracts played by consumers $(k', l') \neq (k, l)$ are still in the menu. Following this recommendation is indeed optimal for type k, l as the new contract is constructed to keep her utility the same and there are no additional deviation strategies. The same holds for all other types with susceptibility k.

For all other types, note that we only have to make sure there are no new deviations into the contract $(x'_{k,l}, R'_{k,l})$ as $\bar{x}$ stays identical. The extremal principle (Proposition 3.4) generalizes to the case with mixed strategies: Suppose there is a contract (x, R) in the support of any type with susceptibility k. A type with susceptibility $k' > k$ acts optimally in choosing (x, R) only if there is no other contract (y, S) chosen by types with susceptibility k with higher consumption $y > x$. Similarly, a contract is only invaded from below ($k' < k$) if it is the contract with the lowest consumption chosen by types with susceptibility k. Since there was no incentive to deviate to the contracts on the edge of the support of $\sigma_{k,l}$, there can be no incentive to deviate to a contract strictly between them giving the same utility to type k, l.

To see this concretely, fix a type k', l' with $k' > k$ and strategy $\sigma^\star_{k',l'}$ and let $y = \max_x \text{support}\sigma_{k,l}$. Then

$$\begin{aligned} u_{k',l'}(\sigma^\star_{k',l'}) \geq u_{k',l'}(y) &= u_{k,l}(y) + (k'-k)\gamma\bar{x}y \\ &= u_{k,l}(\sigma^\star_{k,l}) + (k'-k)\gamma\bar{x}y \\ &= u_{k,l}(x'_{k,l}, R'_{k,l}) + (k'-k)\gamma\bar{x}y \\ &> u_{k,l}(x'_{k,l}, R'_{k,l}) + (k'-k)\gamma\bar{x}x'_{k,l} = u_{k',l'}(x'_{k,l}, R'_{k,l}) \end{aligned}$$

as we needed to show. The proof for $k' < k$ is analogous.

Hence we constructed a new equilibrium, with one type fewer non-degenerately mixing and strictly higher expected revenue for the monopolist. Consequently, for any equilibrium with mixing, there is a strictly better menu of contracts and equilibrium from the point of view of the principal, establishing the claim about revenue maximization. The claim generalizes to mixed strategies with infinite support by an approximation argument. The claim about welfare maximization follows since consumer utility remained unchanged by the above modification to the set of contracts but revenue strictly increased.

For existence, note that revenue is bounded by total surplus, which is given by a quadratic form. Hence, it is sufficient to show that this quadratic form is negative definite.

$$\begin{aligned} R &\leq \sum f_{k,l}\left[(1+\gamma k\bar{x})\, x_{k,l} - \frac{1}{2}x_{k,l}^2\right] \\ &= \sum f_{k,l}\left[\left(1+\gamma k\left(\sum f_{k',l'}\frac{l'}{\mathbb{E}l}x_{k',l'}\right)\right) x_{k,l} - \frac{1}{2}x_{k,l}^2\right] \end{aligned}$$

The symmetric representation of the associated quadratic form is

$$\begin{pmatrix} f_{k,l}^2\gamma\frac{kl}{\mathbb{E}l} & f_{k,l}f_{k',l'}\gamma\frac{kl+k'l'}{\mathbb{E}l} & \cdots \\ & \ddots & \\ & & \end{pmatrix} - \frac{1}{2}\cdot I$$

Note that the maximal eigenvalue of this sum is the maximal eigenvalue of the first matrix minus $\frac{1}{2}$. This matrix is positive and hence, by the Collatz-Wielandt formula, the maximal eigenvalue is bounded by the maximal row sum of the matrix. It is easy to see that the row sum corresponding to type k, l is

$$\frac{1}{2}f_{k,l}\frac{\gamma}{\mathbb{E}l}\left(kl + \mathbb{E}[kl]\right)$$

Hence, the maximum eigenvalue is smaller than $\frac{1}{2}$ if

$$\max_{k,l} \frac{1}{2} f_{k,l} \frac{\gamma}{\mathbb{E}l} \left(kl + \mathbb{E}[kl]\right) < \frac{1}{2}$$

which is satisfied if $\frac{\gamma}{\mathbb{E}l} \left(KL + \mathbb{E}[kl]\right) < 1$.

$\square$

Proof of Proposition 3.1: The first order conditions are immedate. Plugging the expression of $x_{k,l}$ (3.3) into the FOCs of $\bar{x}$ and ζ, we arrive at

$$\zeta = \sum_{k,l} f_{kl} \gamma k \left(1 + \gamma \bar{x} k + \frac{\widetilde{f}_{k,l}}{f_{kl}} \zeta\right) = \gamma \mathbb{E}[k] + \gamma^2 \mathbb{E}[k^2] \bar{x} + \zeta \gamma \frac{\mathbb{E}[kl]}{\mathbb{E}[l]}$$

$$\bar{x} = \sum_{k,l} \tilde{f}_{kl} \left(1 + \gamma \bar{x} k + \frac{\widetilde{f}_{k,l}}{f_{kl}} \zeta\right) = 1 + \gamma \bar{x} \frac{\mathbb{E}[kl]}{\mathbb{E}[l]} + \zeta \frac{\mathbb{E}[l^2]}{\mathbb{E}[l]^2}$$

Solving this 2x2 linear system gives the expressions.

$\square$

Proof of Lemma 3.2: Recall that the relevant set of constraints for this problem are given by $\bigcup_{k \in \mathcal{K}} (k \times \mathcal{L})$. Consider the first-best allocation. The participation constraints are satisfied and the $u_{k,l}$ is independent of l. Hence, by Lemma (3.2), the sorting constraints of this problem are satisfied. Clearly, this is the maximal profit the principal can achieve and hence the first-best allocation is the optimal menu of contracts.

$\square$

Proof of Proposition 3.3: The first point follows immediately from Lemma 3.2. To see the second part, consider the constraints $\mathrm{IC}_{k,l \to k',l'}$ and $\mathrm{IC}_{k',l' \to k,l}$:

$$\left(1 + \gamma \overline{x} k\right) x_{k,l} - \frac{1}{2} x_{k,l}^2 - p_{k,l} \geq \left(1 + \gamma \overline{x} k\right) x_{k',l'} - \frac{1}{2} x_{k',l'}^2 - p_{k',l'}$$

$$\left(1 + \gamma \overline{x} k'\right) x_{k',l'} - \frac{1}{2} x_{k',l'}^2 - p_{k',l'} \geq \left(1 + \gamma \overline{x} k'\right) x_{k,l} - \frac{1}{2} x_{k,l}^2 - p_{k,l}$$

Sum and simplify to arrive at

$$k\left(x_{k,l}-x_{k',l'}\right)\geq k'\left(x_{k,l}-x_{k',l'}\right)$$

as desired.

For part 3

$$\begin{aligned}(1+\gamma\overline{x}k)\,x_{k,l}-\frac{1}{2}x_{k,l}^{2}-p_{k,l} &\geq (1+\gamma\overline{x}k)\,x_{\underline{k},\underline{l}}-\frac{1}{2}x_{\underline{k},\underline{l}}^{2}-p_{\underline{k},\underline{l}}\\ &\geq (1+\gamma\overline{x}\underline{k})\,x_{\underline{k},\underline{l}}-\frac{1}{2}x_{\underline{k},\underline{l}}^{2}-p_{\underline{k},\underline{l}}\\ &\geq 0\end{aligned}$$

□

Proof of Proposition 3.4 (Upward Deviations): Consider k,l and k',l' with $k'>k$. As above

$$\begin{aligned}u_{kl}=u_{kl^{k}} &\geq u_{k+1l_{k+1}}-\gamma\bar{x}x_{k+1,l_{k+1}}\\ &=u_{k+1l^{k+1}}-\gamma\bar{x}x_{k+1,l_{k+1}}\end{aligned}$$

Repeating the above argument, we arrive at

$$\begin{aligned}u_{kl} &\geq u_{k'l^{k'}}-\gamma\bar{x}\sum_{i=k+1}^{k'}x_{il_i}\\ &=u_{k'l'}-\gamma\bar{x}\sum_{i=k+1}^{k'}x_{il_i}\\ &\geq u_{k'l'}-\gamma\bar{x}\sum_{i=k'}^{k-1}x_{k'l'}\\ &=(1+\gamma\overline{x}k)\,x_{k',l'}-\frac{1}{2}x_{k',l'}^{2}-R_{k',l'}\end{aligned}$$

where again the last inequality follows from k-monotonicity and the definition of l_i.

□

Proof of Theorem 3.1: [Rewrite this in utility space immediately, then simplify out the second lemma] We first formulate a reduced problem. In the reduced problem, we keep only local (in k) downward and upward sorting constraints emanating from types with susceptibility L and the participation constraint of type $(\underline{k}, \underline{l})$.

$$\max_{\{p_{k,l},x_{k,l}\}_{k,l\in\mathcal{K}\times\mathcal{L}}} \sum f_{k,l}p_{k,l} \tag{RP}$$

$$\text{s.t.} \qquad \bar{x} = \sum_{k,l} \widetilde{f}_{k,l}x_{k,l} \tag{3.21}$$

$$\forall k, l \in \{\underline{l}, \ldots, L-1\}: \quad (1+\gamma\overline{x}k)\,x_{k,l} - \frac{1}{2}x_{k,l}^2 - p_{k,l} = (1+\gamma\overline{x}k)\,x_{k,l+1} - \frac{1}{2}x_{k,l+1}^2 - R_{k,l+1} \tag{$H_{k,l}$: $\mu_{k,l}$}$$

$$\forall k, l: \quad (1+\gamma\overline{x}k)\,x_{k,L} - \frac{1}{2}x_{k,L}^2 - p_{k,l} \geq (1+\gamma\overline{x}k)\,x_{k-1,l} - \frac{1}{2}x_{k-1,l}^2 - p_{k-1,l} \tag{IC$_{k,L\to k-1,l}$: $\lambda_{k,l}$}$$

$$\forall k, l: \quad (1+\gamma\overline{x}k)\,x_{k,L} - \frac{1}{2}x_{k,L}^2 - p_{k,l} \geq (1+\gamma\overline{x}k)\,x_{k+1,l} - \frac{1}{2}x_{k+1,l}^2 - R_{k+1,l} \tag{IC$_{k,L\to k+1,l}$: $\eta_{k,l}$}$$

$$\forall k, l: \quad (1+\gamma\overline{x}k)\,x_{\underline{k},\underline{l}} - \frac{1}{2}x_{\underline{k},\underline{l}}^2 - R_{\underline{k},\underline{l}} \geq 0 \tag{P:β}$$

□

Lemma. *A menu of contracts solves the reduced problem if and only if it solves the general problem.*

Proof of Lemma. By the arguments given above, the set of horizontal equality constraints ($H_{k,l}$) together with the participation constraint and a full set of extremal constraints is necessary and sufficient. As we impose a set of constraints that is guaranteed to contain an extremal set, this result follows. Note that the level of influence at the source of the sorting constraints is immaterial as utility is constant in L.

Manipulating the Lagrangian of this problem allows us to determine the structure of binding IC constraints and bunching. We use the following Lagrange multipliers and notation conventions:

□

β participation constraint of type $0, 0$

$\mu_{k,l}$	equality constraint for utility between type k, l and $k, l+1$. Hence $\mu_{k,-1} = \mu_{k,L} = 0$.
$\lambda_{k,l}$	sorting constraint from type k, L to type $k-1, l$. Hence $\lambda_{0,l} = \lambda_{K+1,l} = 0$.
λ_k	$\sum_l \lambda_{k,l}$
$\eta_{k,l}$	sorting constraint form type k, L to type $k+1, l$. Hence $\eta_{-1,l} = \eta_{K,l} = 0$.
η_k	$\sum_l \eta_{k,l}$
σ_l	non-negativity constraint on $x_{\underline{k},l}$

Proof. For indices of multiple sums, we write:

$$\sum_{k_i^I l_j^J} = \sum_{k=i}^{I} \sum_{l-j}^{J}$$

and similarly for indices of variables

$$f_{k,l_j^J} = \sum_{l=j}^{J} f_{k,j}.$$

All that being said, we can write the Lagrangian of the problem.

$$\begin{aligned}
\mathcal{L}_{RP} = & \sum_{k,l} p_{k,l} f_{k,l} + \zeta \left(\sum_{k,l} \widetilde{f}_{k,l} x_{k,l} - \bar{x} \right) + \beta \left(x_{\underline{k},\underline{l}} - \frac{1}{2} x^2_{\underline{k},\underline{l}} - R_{\underline{k},\underline{l}} \right) + \\
& + \sum_{k,\underline{l}} \mu_{k,l} \left[(1 + \gamma \overline{x} k) \left(x_{k,l} - x_{k,l+1} \right) - \frac{1}{2} \left(x^2_{k,l} - x^2_{k,l+1} \right) - p_{k,l} + R_{k,l+1} \right] \\
& + \sum_{k_1^K,\underline{l}} \lambda_{k,l} \left[(1 + \gamma \overline{x} k) \left(x_{k,L} - x_{k-1,l} \right) - \frac{1}{2} \left(x^2_{k,L} - x^2_{k-1,l} \right) - p_{k,l} + p_{k-1,l} \right] \\
& + \sum_{\underline{k},\underline{l}} \eta_{k,l} \left[(1 + \gamma \overline{x} k) \left(x_{k,L} - x_{k+1,l} \right) - \frac{1}{2} \left(x^2_{k,L} - x^2_{k+1,l} \right) - p_{k,l} + R_{k+1,l} \right] \\
& + \sum_{l} \sigma_l x_{\underline{k},l}
\end{aligned}$$

Differentiating with respect to $p_{k,l}$ yields

$$f_{k,l}-\chi_{k,l=0,0}\beta-\chi_{l<L}\mu_{k,l}+\chi_{l>\underline{l}}\mu_{k,l-1}+\chi_{k<K}\lambda_{k+1,l}+\chi_{k>\underline{k}}\eta_{k-1,l}-\chi_{l=L}\left[\chi_{k>\underline{k}}\lambda_k+\chi_{k<K}\eta_k\right]=0 \tag{3.22}$$

where χ denotes the indicator function. Summing (3.22) over l, all μ drop out and we get

$$f_k-\chi_{k=0}\beta+\chi_{k<K}\lambda_{k+1}+\chi_{k>\underline{k}}\eta_{k-1}-\left[\chi_{k>\underline{k}}\lambda_k+\chi_{k<K}\eta_k\right]=0$$

Summing this over k, the λ, η drop out and we arrive at

$$\beta=1$$

Returning to the sum above and setting $k=0$, we get

$$\begin{aligned} f_0-1+\lambda_1-\eta_0&=0\\ \lambda_1-\eta_0&=1-f_0 \end{aligned}$$

and the recursion

$$\lambda_{k+1}-\eta_k=\lambda_k-\eta_{k-1}-f_k$$

which we solve for

$$\lambda_{k+1}-\eta_k=1-f_{i_0^k}$$

In the following, we will need to consider which constraints can be binding for some type. By the structure of the problem, the downward constraint $\lambda_{k+1,l}$ is only binding if $x_{k,l}$ is maximal among the $x_{k,\tilde{l}}$ and similarly for $\eta_{k-1,l}$ with minimality. Hence the constraint can only bind at multiple $x_{k,l}$ if they are the same. Furthermore, if at one $x_{k,l}$ both constraints are binding, it is both minimal and maximal and hence all the $x_{k,l}$ for this given k are the same. Similarly, the non-negativity constraints associated with σ_l can only be binding at multiple $x_{\underline{k},l}$ if they are the same and equal to zero.

We now turn to the derivative wrt x_{kl}

$$\begin{aligned} \frac{\partial \mathcal{L}_{RP}}{\partial x_{k,l}} &= \zeta\widetilde{f}_{k,l}+\chi_{k,l=\underline{k},\underline{l}}\beta\left[1+\gamma\bar{x}\underline{k}-x_{\underline{kl}}\right]+\chi_{k=\underline{k}}\sigma_l+\chi_{l<L}\mu_{k,l}\left(1+\gamma\overline{x}k-x_{k,l}\right)\\ &-\chi_{l>\underline{l}}\mu_{k,l-1}\left(1+\gamma\overline{x}k-x_{k,l}\right)-\chi_{l>\underline{l}}\mu_{k,l-1}\left(1+\gamma\overline{x}k-x_{k,l}\right)-\left(\chi_{k<K}\lambda_{k+1,l}+\chi_{k>\underline{k}}\eta_{k-1,l}\right)\left(1+\gamma\overline{x}k-x_{k,l}\right)\\ &-\gamma\bar{x}\left(\chi_{k<K}\lambda_{k+1,l}-\chi_{k>\underline{k}}\eta_{k-1,l}\right)+\chi_{l=L}\left[\chi_{k>\underline{k}}\lambda_k+\chi_{k<K}\eta_k\right]\left(1+\gamma\overline{x}k-x_{k,L}\right)=0 \end{aligned}$$

Plugging the multipliers from $\frac{\partial \mathcal{L}_{RP}}{\partial p_{k,l}}$ into the above we get

$$0 = \zeta \widetilde{f}_{k,l} + f_{k,l}\left[1 + \gamma \overline{x} k - x_{k,l}\right] + \chi_{k=\underline{k}} \sigma_l - \gamma \bar{x}\left(\chi_{k<K} \lambda_{k+1,l} - \chi_{k>\underline{k}} \eta_{k-1,l}\right)$$

and hence

$$\begin{aligned} x_{k,l} &= 1 + \gamma \bar{x} k + \zeta \frac{l}{\mathbb{E}[l]} + \chi_{k=\underline{k}} \sigma_l - \gamma \bar{x}\left(\chi_{k<K} \lambda_{k+1,l} - \chi_{k>\underline{k}} \eta_{k-1,l}\right) \\ &= \check{x_{k,l}} + \chi_{k=\underline{k}} \sigma_l - \gamma \bar{x}\left(\chi_{k<K} \lambda_{k+1,l} - \chi_{k>\underline{k}} \eta_{k-1,l}\right) \end{aligned}$$

where $\check{x_{k,l}}$ is the efficient level of consumption for type k, l given $\bar{x}, \zeta$.

Suppose that $\zeta \geq 0$ and there are a k, l, l' with $l < l'$ and $x_{k,l} > x_{k,l'}$. We know that $\eta_{k-1,l} = \lambda_{k+1,l'} = 0$ since the respective x are not extremal. Furthermore, if $k = \underline{k}$, we have $\sigma_l = 0$. Hence we get

$$\begin{aligned} x_{k,l} &= \check{x_{k,l}} - \gamma \bar{x} \chi_{k<K} \lambda_{k+1,l} \\ &\leq \check{x_{k,l}} \\ &\leq \check{x_{k,l'}} \\ &\leq \check{x_{k,l'}} + \chi_{k=\underline{k}} \sigma_l + \gamma \bar{x} \chi_{k>\underline{k}} \eta_{k-1,l} \\ &= x_{k,l'} < x_{k,l} \end{aligned}$$

a contradiction. The proof for $\zeta \leq 0$ is analogous. It remains to show that $\zeta > 0$.

□

Lemma. *$\zeta > 0$.*

Proof of Lemma: Let us first formulate the utility space problem with increasing

$$\max_{\{u_k\}_{k \in \mathcal{K}}, \{x_{k,l}\}_{k,l \in \mathcal{K} \times \mathcal{L}}} \sum f_{k,l}\left[\left(1 + \gamma \overline{x} k\right) x_{k,l} - 0.5 x_{k,l}^2 - u_k\right] \tag{UP}$$

$$\text{s.t.} \quad \bar{x} = \sum_{k,l} \widetilde{f}_{k,l} x_{k,l} \tag{3.23}$$

$$\forall k \in \{\underline{k}+1, \ldots, K\}: \quad u_k = u_{k-1} + \gamma \bar{x} x_{k-1,L} \qquad (\text{IC}_{k \to k-1}\colon \lambda_k)$$

$$\forall k, l \in \mathcal{K} \times \{\underline{l}+1, \ldots, L\}: \quad x_{k,l} \geq x_{k,l-1} l > \underline{l} \qquad (\ \sigma_{k,l})$$

$$\forall k \in \{\underline{k}+1, \ldots, K\}: \quad x_{k,\underline{l}} \geq x_{k-1,L} \qquad (:\sigma_{k,l})$$

$$x_{\underline{k},\underline{l} \geq 0} \qquad (:\sigma_{\underline{k},\underline{l}})$$

$$u_{\underline{k}} \geq 0 \qquad (\text{P}:\beta)$$

and with reverse order

$$\max_{\{u_k\}_{k\in\mathcal{K}},\{x_{k,l}\}_{k,l\in\mathcal{K}\times\mathcal{L}}} \sum f_{k,l}\left[\left(1+\gamma\bar{x}k\right)x_{k,l}-0.5x_{k,l}^2-u_k\right] \tag{UP'}$$

$$\text{s.t.} \quad \bar{x}=\sum_{k,l}\widetilde{f}_{k,l}x_{k,l} \tag{3.24}$$

$$\forall k\in\{\underline{k}+1,\ldots,K\}: \quad u_k=u_{k-1}+\gamma\bar{x}x_{k-1,\underline{l}} \qquad (\mathrm{I}C_{k\to k-1}\colon \lambda_k)$$

$$\forall k,l\in\mathcal{K}\times\{\underline{l},\ldots,L-1\}: \quad x_{k,l}\geq x_{k,l+1}l>\underline{l} \qquad (\ \sigma_{k,l})$$

$$\forall k\in\{\underline{k}+1,\ldots,K\}: \quad x_{k,L}\geq x_{k-1,\underline{l}} \qquad (\colon \sigma_{k,L})$$

$$x_{\underline{k},L\geq 0} \qquad (\colon \sigma_{\underline{k},L})$$

$$u_{\underline{k}}\geq 0 \qquad (\mathrm{P}{:}\beta)$$

To see that a feasible allocation in either UP or UP' is feasible in GP, all we need to show is that it is feasible in RP. Equality of utility for fixed k across l is enforced directly, as u_k is independent of l. From lexicographic monotonicity, the constraint $(k,\underline{l})\to(k-1,L)$ or $(k,L)\to(k-1,\underline{l})$ is an extremal downward constraint for all k. To show upward sorting, note that for the deviation payoff from $(k-1,L)$ to $(k,\underline{l})$ (analogously for the case $\zeta<0$ from $(k-1,\underline{l})$ to (k,L)) we have

$$\begin{aligned}
(1+\gamma\bar{x}(k-1))x_{k,\underline{l}}-\frac{1}{2}x_{k,l_o}^2-R_{k,\underline{l}} &= (1+\gamma\bar{x}k)x_{k,\underline{l}}-\frac{1}{2}x_{k,l_o}^2-R_{k,\underline{l}}-\gamma\bar{x}x_{k,\underline{l}}\\
&\leq (1+\gamma\bar{x}k)x_{k-1,L}-\frac{1}{2}x_{k-1,L}^2-p_{k-1,l}-\gamma\bar{x}x_{k,\underline{l}}\\
&= (1+\gamma\bar{x}(k-1))x_{k-1,L}-\frac{1}{2}x_{k-1,L}^2-p_{k-1,l}-\gamma\bar{x}x_{k,\underline{l}}+\gamma\bar{x}x_{k-1,L}\\
&= u_{k-1,L}-\gamma\bar{x}\left(x_{k,\underline{l}}-x_{k-1,L}\right)\leq u_{k-1,L}
\end{aligned}$$

and the upward sorting constraint is satisfied as well. Hence we have a full set of extremal sorting constraints and by 3.4 the allocation satisfies all sorting constraints. All participation constraints are implied by the participation constraint of the lowest type, as are all non-negativity constraints by the non-negativity of $x_{\underline{k},\underline{l}}$and lexicographic monotonicity.

To show that the solution to GP is the solution either to UP or UP', note that in the solution to the GP, either $\zeta\geq 0$ or $\zeta<0$. In the former case, the solution satisfies lexicographic monotonicity by Proposition 3.1 and is feasible and optimal in UP. In the latter case, it satisfies lexicographic monotonicity in $\mathcal{K}\times(-\mathcal{L})$ and is feasible and optimal in UP'.

The downward sorting constraints are binding since $\lambda_k > 0$ for all k, so at least one of the (k, l) for given k is constrained, which has to be (k, L) by lexicographic monotonicity and extremal sorting.

The proof for the claims for $\zeta < 0$ are identical, exchanging $\underline{l}$ and L where applicable.

Computations for the utility-space problem

The Lagrangian for the problem is given by

$$\begin{aligned}\mathcal{L}_{UP} = & \sum f_{k,l}\left[\left(1+\gamma\overline{x}k\right)x_{k,l} - \frac{1}{2}x_{k,l}^2 - u_k\right] + \zeta\left(\sum_{k,l}\widetilde{f}_{k,l}x_{k,l} - \bar{x}\right) + \beta u_0 + \\ & + \sum_{k_1^K}\lambda_k\left[u_k - u_{k-1} - \gamma\bar{x}x_{k-1,L}\right] \\ & + \sum_{k_{\underline{k}}^K, l_{\underline{l}+1}^L}\sigma_{k,l}\left[x_{k,l} - x_{k,l-1}\right] \\ & + \sum_{k_{\underline{k}+1}^K}\sigma_{k,\underline{l}}\left[x_{k,\underline{l}} - x_{k-1,L}\right] + \sigma_{\underline{k},\underline{l}}\left[x_{\underline{k},\underline{l}}\right]\end{aligned}$$

Taking first order conditions:

$$\frac{\partial\mathcal{L}_{UP}}{\partial x_{k,l}} = f_{k,l}\left[\left(1+\gamma\overline{x}k\right) - x_{k,l}\right] + \zeta\widetilde{f}_{k,l} + \sigma_{k,l} - \chi_{l\neq L}\sigma_{k,l+1} - \chi_{k<K,l=L}\left(\gamma\overline{x}\lambda_{k+1} + \sigma_{k+1,\underline{l}}\right) = 0$$

$$\frac{\partial\mathcal{L}_{UP}}{\partial u_k} = -f_k + \chi_{k>\underline{k}}\lambda_k - \chi_{k<K}\lambda_{k+1} + \chi_{k=0}\beta = 0$$

Summing over all $\frac{\partial\mathcal{L}_{UP}}{\partial u_k}$

$$\sum_k \frac{\partial\mathcal{L}_{UP}}{\partial u_k} = 1 - \beta = 0$$

we get $\beta = 1$. Furthermore

$$\begin{aligned}\lambda_{\underline{k}+1} &= 1 - f_{\underline{k}} \\ \lambda_{k+1} &= \lambda_k - f_k \\ \lambda_{k+1} &= 1 - F_k\end{aligned}$$

Finally, consider

$$\frac{\partial \mathcal{L}_{UP}}{\partial \bar{x}} = \sum f_{k,l}\gamma k x_{k,l} - \zeta - \sum_{k_1^K} \lambda_k \gamma x_{k-1,L} = 0$$

$$\begin{aligned}
\zeta &= \sum f_{k,l}\gamma k x_{k,l} - \sum_{k_1^K} \lambda_k \gamma x_{k-1,L} \\
&= \gamma \left\{ \sum f_{k,l} k x_{k,l} - \sum_{\underline{k}}^{K-1} (1 - F_k)\, x_{k,L} \right\} \\
&= \gamma \left\{ \sum_{\underline{k}}^{K} [k f_k \mathbb{E}[x_{k,l}|k] - (1 - F_k)\, x_{k,L}] \right\}
\end{aligned}$$

Let $x_k := \mathbb{E}[x_{k,l}|k]$ and note that

$$\begin{aligned}
\sum_{\underline{k}}^{K} k f_k x_k &= \sum_{\underline{k}}^{K} (F_k - F_{k-1})\, k x_k = \sum_{\underline{k}}^{K} F_k k x_k - \sum_{\underline{k}}^{K} F_{k-1} k x_k \\
&= \sum_{\underline{k}}^{K} F_k k x_k - \sum_{\underline{k}-1}^{K-1} F_k\, (k+1)\, x_{k+1} \\
&= \sum_{\underline{k}}^{K} F_k k x_k - \left[\sum_{\underline{k}}^{K} F_k\, (k+1)\, x_{k+1} + \underbrace{F_{\underline{k}-1} \underline{k} x_{\underline{k}}}_{=0} - \underbrace{F_K}_{=1} (K+1)\, x_{K+1} \right] \\
&= (K+1)\, x_{K+1} - \sum_{\underline{k}}^{K} F_k\, [(k+1)\, x_{k+1} - k x_k] \\
&= \sum_{\underline{k}}^{K} (1 - F_k)\, [(k+1)\, x_{k+1} - k x_k]
\end{aligned}$$

Note that x_{K+1} is an arbitrary extension of the sequence and will cancel out of the expression. Plugging back, we get

$$\begin{aligned}\zeta &= \gamma\left\{\sum_{\underline{k}}^{K}(1-F_k)\left[(k+1)\,x_{k+1}-kx_k\right]-(1-F_k)\,x_{k,L}\right\}\\ &= \gamma\left\{\sum_{\underline{k}}^{K}(1-F_k)\left[kx_{k+1}-kx_k\right]+(1-F_k)\left[x_{k+1}-x_{k,L}\right]\right\}\\ &\geq \gamma\left\{\sum_{\underline{k}}^{K}(1-F_k)\,k\left[x_{k+1}-x_k\right]\right\}\geq 0\end{aligned}$$

with strict inequality, when there isn't complete pooling. An analogous argument works for (UP'). Hence $\zeta \geq 0$ and the original problem satisfies the lexicographic order on $\mathcal{K}\times\mathcal{L}$.

Which concludes the proof of the Theorem.

□

Proof of Proposition 3.5: This follows immediately from the utility-space problem above, noticing that the participation constraint has to bind and summing the utility term by parts.

□

Returning to x, let us define a block b of indices as a set of indices that satisfies the following three conditions

1. if $(k,l)\in b \& (k',l')\in b$, then $(k'',l'')\in b$ for all (k'',l'') satisfying $(k,l)<(k'',l'')<(k',l')$ (i.e. b is an order interval in the lexicographic order underlying the problem),

2. if $(k,l)\in b$, then either $\sigma_{k,l}>0$, $(k,l)=\min b$, and

3. b is maximal with respect to these properties.

In other words, we call a set of indices a block, if it is an interval of indices between which the monotonicity constraint binds. Note that a single index $\{(k,l)\}$ is a block according to this condition iff both the downward ($\sigma_{k,l}$) and upward ($\sigma_{k,l+1}$) monotonicity constraints are slack, or it is the minimal index ($\underline{k}, \underline{l}$) (in case the non-negativity constraint is binding). It is easy to see that the set of blocks forms a partition of $\mathcal{K} \times \mathcal{L}$ and that $x_{k,l}$ is measurable with respect to this partition. We can order the blocks by extending the lexicographic order to them.

Denote x_b the level of consumption on the block b, similarly $f_b = \sum_{(k,l)\in b} f_{k,l}$. If $(\underline{k}, \underline{l}) \in b$ and $\sigma_{\underline{k},\underline{l}} > 0$, we get $x_b = 0$. Otherwise, let us sum the FOC wrt x over a block to arrive at

$$\sum_b \frac{\partial \mathcal{L}_{UP}}{\partial x_{k,l}} = f_b \left[(1 + \gamma \overline{x} \mathbb{E}[k|b]) - x_b \right] + \zeta \widetilde{f}_b + \sigma_{\min b} - \sigma_{\max b+1} - \sum_{\{k:(k,L)\in b\}} \gamma \overline{x} \lambda_{k+1} = 0$$

As by definition of a block, $\sigma_{\min b} = 0$, except in the case of the lowest block, which we treated separately. Similarly, for every block $\sigma_{\max b+1} = 0$, hence

$$x_b = 1 + \gamma \overline{x} \mathbb{E}[k|b] + \zeta \frac{\mathbb{E}[l|b]}{\mathbb{E}[l]} - \gamma \bar{x} \frac{\sum_{\{k:(k,L)\in b\}} 1 - F_k}{f_b}$$

Structure of Monotonicity Blocks

Note that because the only "sources of non-monotonicity" are adjacent to k, L, every bunching region needs to contain such a type.

Proof of Lemma 3.4: Conditional on a block structure, the x_{kl} are continous. But note that the convexification procedure in our case can be rewritten as a continous operations on the finite-dimensional space of virtual values that is continous in virtual values. Hence, Γ is continous.

Note that

$$\begin{aligned}
\zeta &= \gamma \left\{ \sum f_{k,l} k x_{k,l} - \sum_{\underline{k}}^{K-1} (1 - F_k) \, x_{k,L} \right\} + \kappa c'(\bar{x}) \\
&\leq \gamma \sum f_{k,l} k x_{k,l} + \kappa \leq \gamma \sum f_{k,l} k x_{K,L} + \kappa = \gamma \mathbb{E}[k] x_{K,L} + \kappa \\
&\leq \gamma \mathbb{E}[k] \left(1 + \gamma K \bar{x} + \zeta \frac{L}{\mathbb{E}[l]} \right) + \kappa
\end{aligned}$$

To bound aggregate consumption

$$\begin{aligned}\bar{x} &= \sum_{k,l} \widetilde{f}_{k,l} x_{k,l} \\ &\leq \frac{L}{\mathbb{E}[l]} x_{K,L} \\ &\leq \frac{L}{\mathbb{E}[l]} \left(1 + \gamma K \bar{x} + \zeta \frac{L}{\mathbb{E}[l]}\right)\end{aligned}$$

We arrive at

$$\begin{pmatrix} \bar{x} \\ \zeta \end{pmatrix} \leq \begin{pmatrix} \frac{L}{\mathbb{E}[l]} \\ \gamma \mathbb{E}[k] + \kappa \end{pmatrix} + \begin{pmatrix} \frac{L}{\mathbb{E}[l]} \gamma K & \frac{L^2}{\mathbb{E}[l]^2} \\ \gamma \mathbb{E}[k] \gamma K & \gamma \mathbb{E}[k] \frac{L}{\mathbb{E}[l]} \end{pmatrix} \begin{pmatrix} \bar{x} \\ \zeta \end{pmatrix}$$

$$\begin{pmatrix} 1 - \frac{L}{\mathbb{E}[l]} \gamma K & \frac{L^2}{\mathbb{E}[l]^2} \\ \gamma \mathbb{E}[k] \gamma K & 1 - \gamma \mathbb{E}[k] \frac{L}{\mathbb{E}[l]} \end{pmatrix} \begin{pmatrix} \bar{x} \\ \zeta \end{pmatrix} \leq \begin{pmatrix} \frac{L}{\mathbb{E}[l]} \\ \gamma \mathbb{E}[k] + \kappa \end{pmatrix}$$

The matrix is positive definite if $1 - \frac{L}{\mathbb{E}[l]} \gamma K - \gamma \mathbb{E}[k] \frac{L}{\mathbb{E}[l]} > 0$, i.e. $1 > \gamma \frac{L}{\mathbb{E}[l]} \left(K + \mathbb{E}[k]\right)$. Then, we get upper bounds by inverting the matrix.

□

Proof of Theorem (3.2): Consider Γ as defined on co(cl(range(Γ))). This set is convex and compact, as the range is precompact and the convex hull of compact sets is compact in finite-dimensional vector spaces. The existence of a fixed point follows from Brouwer's theorem. We established that the optimal contract takes the desired form for any given $\bar{x}, \zeta$. Equation (3.14) encodes consistency conditions. The first component is the definition of $\bar{x}$ and therefore required by feasibility. The second component is the definition of ζ, the proof of optimality conditional on $\bar{x}, \zeta$ is predicated on this equation being satisfied. Hence, both are necessary conditions for an optimum.

□

Proof of Proposition 3.6: We can rewrite the problem in utility space, noting that $u_{kl} = (1 + \gamma \overline{x} k)\, x_{k,l} - \frac{1}{2} x_{k,l}^2 - p_{k,l}$ or equivalently $p_{kl} = (1 + \gamma \overline{x} k)\, x_{k,l} - \frac{1}{2} x_{k,l}^2 - u_{kl}$. Then, P is equivalent to $u_{k_0 l} = 0$ where equality follows from by the usual argument. IC is equivalent to $u_{k,l} \geq u_{k'l} + \gamma \overline{x} \left(k - k'\right) x_{k',l}$. Again, by the usual arguments, local downward IC and monotonicity are sufficient and IC are binding, hence $u_{k,l} =$

$\gamma\bar{x}\sum_{i=k_0}^{k-1} x_{i,l}$. Plugging this into the objective and applying summation by parts to the double sum, we arrive at the Proposition.

□

Proof of Proposition 3.7: Let $F_k^l = \sum_{i=\underline{k}}^{k} f_{i,l}$. Then $\zeta = \gamma \sum_{k,l} f_{kl} \left\{ k - \frac{F_K^l - F_k^l}{f_{kl}} \right\} x_{kl}$ and the claim about ζ follows since

$$\begin{aligned}
\sum_{k,l} k f_{kl} x_{kl} = \sum_l \sum_k k f_{kl} x_{kl} = \sum_l \sum_{\underline{k}}^{K} \left(F_k^l - F_{k-1}^l\right) k x_{kl} &= \sum_l \left\{ \sum_{\underline{k}}^{K} F_k^l k x_{kl} - \sum_{\underline{k}}^{K} F_{k-1}^l k x_{kl} \right\} \\
&= \sum_l \left\{ \sum_{\underline{k}}^{K} F_k^l k x_{kl} - \sum_{\underline{k}-1}^{K-1} F_k^l \left(k+1\right) x_{k+1,l} \right\} \\
&= \sum_l \left\{ \sum_{\underline{k}}^{K} F_k^l k x_{kl} - \left[\sum_{\underline{k}}^{K} F_k^l \left(k+1\right) x_{k+1,l} + \underbrace{F_{\underline{k}-1}^l \underline{k} x_{\underline{k},l}}_{=0} - F_K^l \left(K+1\right) x \right.\right. \\
&= \sum_l \left\{ F_K^l \left(K+1\right) x_{K+1,l} - \sum_{\underline{k}}^{K} F_k^l \left[\left(k+1\right) x_{k+1,l} - k x_{k,l}\right] \right\} \\
&= \sum_l \sum_{\underline{k}}^{K} \left(F_K^l - F_k^l\right) \left[\left(k+1\right) x_{k+1,l} - k x_{k,l}\right]
\end{aligned}$$

Note that $x_{K+1,l}$ is an arbitrary extension of the sequence and will cancel out of the expression. Plugging back, we get

$$\begin{aligned}
\zeta &= \gamma \left\{ \sum_l \sum_{\underline{k}}^{K} \left(F_K^l - F_k^l\right) \left[\left(k+1\right) x_{k+1,l} - k x_{k,l}\right] - \left(F_K^l - F_k^l\right) x_{kl} \right\} \\
&= \gamma \left\{ \sum_l \sum_{\underline{k}}^{K} \left(F_K^l - F_k^l\right) \left[k x_{k+1,l} - k x_{k,l}\right] + \left(F_K^l - F_k^l\right) \left[x_{k+1,l} - x_{k,l}\right] \right\} \\
&= \gamma \left\{ \sum_l \sum_{\underline{k}}^{K} \left(F_K^l - F_k^l\right) \left(k+1\right) \left[x_{k+1,l} - x_{k,l}\right] \right\} > 0
\end{aligned}$$

Note that the system is fully linear, so we can solve it (a little bit)

$$
\begin{aligned}
x_{kl} &= 1 + \gamma \bar{x}\left(k - \frac{F_K^l - F_k^l}{f_{kl}}\right) + \frac{l}{\mathbb{E}[l]}\zeta \\
\zeta &= \gamma \sum_{k,l} f_{kl}\left\{k - \frac{F_K^l - F_k^l}{f_{kl}}\right\} x_{kl} \\
\bar{x} &= \sum_{k,l} f_{kl} l \left\{1 + \gamma \bar{x}\left(k - \frac{F_K^l - F_k^l}{f_{kl}}\right) + \frac{l}{\mathbb{E}[l]}\zeta\right\} \\
&= \mathbb{E}[l] + \gamma \bar{x}\left(\mathbb{E}[kl] - \sum_{k,l} f_{kl} l \frac{F_K^l - F_k^l}{f_{kl}}\right) + \frac{\mathbb{E}[l^2]}{\mathbb{E}[l]}\zeta \\
&= \mathbb{E}[l] + \gamma \bar{x}\left(\mathbb{E}[kl] - \sum_{k,l} l\left(F_K^l - F_k^l\right)\right) + \frac{\mathbb{E}[l^2]}{\mathbb{E}[l]}\zeta \\
&= \mathbb{E}[l] + \gamma \bar{x}\left(\underline{k}\mathbb{E}[l]\right) + \frac{\mathbb{E}[l^2]}{\mathbb{E}[l]}\zeta \\
\bar{x} &= \frac{1}{1 - \gamma \underline{k}\mathbb{E}[l]}\left\{\mathbb{E}[l] + \frac{\mathbb{E}[l^2]}{\mathbb{E}[l]}\zeta\right\} \\
\zeta &= \gamma \left\{\sum_{l} \sum_{\underline{k}}^{K} \left(F_K^l - F_k^l\right)(k+1)\gamma \bar{x}\left[k + 1 - \frac{F_K^l - F_{k+1}^l}{f_{k+1,l}} - k + \frac{F_K^l - F_k^l}{f_{kl}}\right]\right\} \\
&= \gamma^2 \bar{x}\left\{\sum_{l} \sum_{\underline{k}}^{K} \left(F_K^l - F_k^l\right)(k+1)\left[1 - \frac{F_K^l - F_{k+1}^l}{f_{k+1,l}} + \frac{F_K^l - F_k^l}{f_{kl}}\right]\right\} \\
\zeta &= \gamma^2 \bar{x}\Xi = \frac{\gamma^2 \Xi}{1 - \gamma\left(\mathbb{E}[kl] - \Psi\right)}\left\{\mathbb{E}[l] + \frac{\mathbb{E}[l^2]}{\mathbb{E}[l]}\zeta\right\}
\end{aligned}
$$

Where we noticed that

$$
\sum_{k,l} l\left(F_K^l - F_k^l\right) = \sum_{k,l} l f_l\left(1 - \frac{F_k^l}{f_l}\right) = \sum_{l} l f_l \sum_{k=\underline{k}}^{K}\left(1 - \frac{F_k^l}{f_l}\right) = \sum_{l} l f_l\left(\mathbb{E}[k|l] - \underline{k}\right) = \mathbb{E}[kl] - \underline{k}\mathbb{E}[l]
$$

The level of consumption is increasing in l if

$$
\frac{l}{\mathbb{E}[l]}\zeta - \gamma \bar{x}\frac{F_K^l - F_k^l}{f_{kl}}
$$

$$
\frac{l}{\mathbb{E}[l]}\gamma \Xi - \frac{F_K^l - F_k^l}{f_{kl}}
$$

is increasing in l.

□