Collective Quality *

Pierre Fleckinger[†] Wanda Mimra[‡] Angelo Zago[§]

October 8, 2024

Abstract

The prevalence of collective names and brands presents a puzzle, given the potential for free-riding and misaligned incentives often associated with collective production. This paper develops a simple model of collective reputation for quality. Producers with privately known costs decide whether to invest in high quality, which is only imperfectly observed by the market. Mechanisms such as expert inspections, certification, or regulatory oversight reveal high-quality products with some probability. However, for unidentified products, market assessments are based on the collective reputation structure, which pools producers into groups. We establish that equilibrium in this setting is unique and highlight a positive incentive effect of free-riding within groups: the quality efforts of more efficient producers are enhanced. Consequently, aggregate quality and welfare can be higher under collective reputations compared to individual ones. Despite these potential gains, groups tend to unravel in the absence of transfers. Nonetheless, we propose intuitive, type-independent, and budget-balanced group contracts that can sustain optimal quality provision. This model of collective reputation has broad applications, including the design of admission thresholds in education.

JEL classification: D82; D47; D71; L15.

Keywords: Collective Reputation, Quality Provision, Incentives, Pooling, Free-riding, Reputation Milking.

*We would like to thank Heski Bar-Isaac, Elchanan Ben-Porath, Arthur Fishman, Alex Gershkov, Paul Heidhues, Yassine Lefouili, Raphaël Levy, Zvika Neeman, Aniko Öry, Takuro Yamashita and seminar participants at Düsseldorf Institute for Competition Economics, INRAe Aliss, Tel Aviv University, Hebrew University, Paris School of Economics, Toulouse School of Economics, University of Montpellier, University of Nancy, University of Verona and University of Zürich for very helpful comments and suggestions.

†Mines Paris-PSL and Paris School of Economics. Email: `pierre.fleckinger@minesparis.psl.eu`.

‡ESCP Business School. Email: `wmimra@escp.eu`.

§University of Verona. Email: `angelo.zago@univr.it`.

Contents

1 Introduction

The prevalence of collective names and brands may appear perplexing, as collective production is typically associated with free-riding and non-tailored incentives. However, collective reputation remains a crucial institution in markets where asymmetric information about product quality prevails, and it frequently proves effective in practice. Champagne, probably the most globally renowned wine, is a prime example. Over 16,000 small vine-growing producers and 300 Champagne houses, collectively shipping more than 300 million bottles annually, all share in the prestige associated with the Champagne label, a brand whose protection was even included in the Treaty of Versailles after World War I.[1]

Collective reputation is not only ubiquitous in agricultural markets; it is quite common in manufacturing industries – Swiss watches, high fashion from Milano, machinery or cars made in Germany – where country of origin is perceived as evocative of the (high) average quality of products. It also plays an important role in many service sectors and regulated professions. The organization of education constitutes another example. A school's reputation affects the choices of students, families, faculty and school administrators[2] and it has been long recognized that graduates use their school of origin to signal ability, endogenously leading to stratification.[3]

Despite its importance, the functioning of quality provision under a collective name remains poorly understood compared to individual reputation mechanisms. In this paper, we propose a simple and tractable model of collective reputation for quality. Our notion of collective reputation is based on two features: (1) heterogeneous producers are pooled into groups, and (2) quality is imperfectly observed but buyers draw inferences based on statistical knowledge available on the group a producer belongs to. Pooling producers leads to some free-riding, but we demonstrate *positive* incentive effects of this free-riding, and derive the associated efficiency gains. In essence, we show that the pooling inherent to a collective organization can both result in higher aggregate quality and lower total costs than what producers could achieve individually. A key advantage of our modeling approach is to replicate dynamic reputation effects in a parsimonious static model. This allows us to introduce a rich type space and fully characterize all equilibria, facilitating comparative statics and paving the way for tractable applications.

[1]In 2015, Champagne hillsides, houses and cellars were immortalized by inclusion in UNESCO's World Heritage list.

[2]See, e.g., Drori et al. (2015).

[3]See MacLeod and Urquiola (2015).

Our theory of collective quality leans on an important characteristic of the information available to buyers: Individual quality is not always observed, and verifiable information on quality tends to exhibit bias, often towards positive news. For instance, schools and universities award degrees to successful students, providing them with positive certification, while instances of student dropout are not publicized. In the wine market, experts often exhibit a favorable bias, tending to report predominantly positive news while omitting or downplaying negative evaluations. A typical illustration is that a wine trades at a higher price when it has received *any* Parker grade, rather than received none.[4] Of course, in other instances, such as those involving corruption or fraud, reputation is either maintained, or compromised when bad news emerge. While feedback provided by professors tends to be positive, news delivered by the police is often negative. The exposition focuses on the first kind of environment, that we term *friendly*, but we also provide a complete discussion of the alternative *neutral* and *hostile* environments.[5]

Our model features a continuum of producers with heterogeneous costs of providing high quality. These producers are organized into disjoint groups under a collective reputation structure. Buyers do not initially observe the true quality of the product. However, the market can detect high quality with some probability through processes such as inspections, certifications, recommendations from experts, award ceremonies and similar practices. Products (or services) not detected as of high quality are pooled together according to the collective reputation structure, i.e., the partitioning of producers. In such a case, buyers are only aware that the producer belongs to a given group, which composition is known. As a result, incentives for producers originate from two different channels: A direct one through the imperfect detection of high quality products, and an indirect one through the market's inference of the quality of unidentified products. The collective dimension kicks in with the indirect incentive effect: The market's belief over the quality of an unidentified product depends on which group its producer belongs to.

An important reference case corresponds to individually identified producers. Reputation is in that case tailored to the characteristic of each given producer–here, their cost of providing high quality. Moral hazard on quality provision remains an issue, but pooling is

[4]This selectively positive bias has been well-documented, see for instance Ali et al. (2008). This bias is explained by a number of factors, such as the presence of incumbent experts protecting their market influence, or experts being paid by producers, or by selective samples submitted by producers for evaluation.

[5]We do not model why a particular informational environment prevails, but take it as given. The seminal analysis of persisting corrupt reputation by Tirole (1996) relies on a given hostile environment. In contrast, we emphasize here friendly environments and borrow the terminology and information structure from Fleckinger et al. (2017). Appendix A discusses in detail the other informational environments.

absent. In a friendly environment, a good reputation is self-defeating because producers are tempted to milk their reputation. Indeed, suppose the market believes that the producer supplies high quality, so that reputation is high. Then incentives to provide high quality and get a premium over the baseline reputation are weaker than they would be under a lower reputation. As a result, in the unique equilibrium, a producer milks his individual reputation to some extent, and will not produce quality at the efficient level.

When reputation is collective, heterogeneous producers are grouped together, and the market does not identify a given producer individually. Similarly to the case of individual reputation, a high reputation hinders incentives. However, since a collective reputation is based on the behavior of many (heterogeneous) agents, it differs in general from individual reputation. So do incentives. We first show that any collective reputation structure yields a unique equilibrium. Equilibrium is characterized by a cost threshold in each group, such that only the members with costs below the threshold provide full quality, while the others fully free-ride. But while in the individual case reputation milking is a pure efficiency loss, free-riding in the group decreases the reputation, which creates in turn incentives for other group members. We show that, as a result, collective reputation can yield higher quality than individual reputation. In addition, the pattern of free-riding under collective reputation leads to an efficient cost allocation for the aggregate quality produced. These effects in combination explain why collective reputation may spur quality and increase welfare.

Unfortunately, despite the possibility of higher aggregate profits (and welfare) under collective reputation, this benefit does not hold at the individual level. Lowest cost producers in a group always prefer to secede and operate under individual reputations: any group with collective reputation without transfers or mandatory participation tends to unravel. Contrary to the case of pure adverse selection emphasized in the seminal model by Milgrom (1981), such unraveling can be inefficient, in which case it should be prevented.

Designing a stable, budget-balanced collective organization with internal transfers proves to be a fruitful exercise. We propose two simple type-independent and budget-balanced transfer schemes based solely on realized market price that implement the first-best efficient outcome. In the first scheme, which we call a "bonus club," producers pay a membership fee upfront, and these funds are later distributed as bonuses in addition to the market prices they receive. A well-calibrated bonus scheme allows to increase incentives to the efficient level. An alternative scheme is a "common retailer" organization, in which producers sell through a collective channel that distorts prices paid back to the producer, so as to generate efficient incentives. The advantage is that producers do not need liquidity initially to implement this second solution. Note that no further information is generated in either arrangements: only the market price is used in the redistribution scheme. Strikingly, both these collective organizations fully leverage the benefits of collective reputation to restore efficient incentives.

The last section applies our theory of collective quality to admission thresholds in education. We show in particular that creating tiered groups always improves on full collective reputation, but that a badly chosen admission threshold can conversely reduce welfare. We believe our model can shed new light to study education systems, and to understand the rationale and properties of public elitism, in countries such as France, India, China and other Asian countries, as well as the functioning of more profit-guided education systems in other countries, such as the US and the UK.

Related literature

While the interplay of individual reputation and quality provision spans a vast literature (see MacLeod, 2007; Bar-Isaac and Tadelis, 2008, for comprehensive surveys), contributions to our understanding of collective reputation are still few, and they take different approaches towards modeling the very notion of collective reputation. Our focus is on the incentives provided, and resulting efficiency, when pooling heterogeneous agents under a collective name. The collective reputation dimension comes into play in the way the market's inference about quality depends on the composition of producer groups. This inference process in turn shapes incentives. To introduce the rich type space required for an analysis based on group composition, we propose a model that succinctly captures key dynamic reputation effects within a parsimonious static framework.

The seminal paper by Tirole (1996) builds on the classic game theoretic approach to reputation (*à la* Kreps and Wilson, 1982), where agents are of different behavioral types and whose actions are partially revealing. Tirole (1996) studies dynamic situations, with a focus on the interplay between individual and collective reputation, as well as on the circumstances under which low quality equilibria may persist, locking organizations in a collective reputational trap.[6] Imperfect observability creates history dependence, and in

[6]Blume (2006) and Levin (2009) consider stochastic versions of Tirole (1996)'s collective reputation model. Levin (2009) shows in particular why moving from one steady-state to another may be gradual, and why small policy changes may fail to favorably shift behavior following a history of poor outcomes. In the closely-related literature on statistical discrimination, Kim and Loury (2018) pursue this line of inquiry by extending the original static setting of Coate and Loury (1993) to a dynamic version. Our main interest

particular, historical low quality indicates corrupt agents. This raises general suspicion and thereby affects new agents that cannot be distinguished from corrupt ones. A key feature of the framework used by Tirole (1996) is the particular informational environment: it is hostile, since the only news that surface consist of bad quality events.[7] Incentives for quality and beliefs then go hand-in-hand: if the market believes quality is low, incentives are low, since no signal can contradict this market belief, which becomes self-fulfilling. If the market believes that quality is high, in turn, incentives to maintain this belief indeed arise from the fear of being detected when cheating. In our model below, we concentrate on the diametrically opposed informational environment in which only high quality can be revealed. In this environment, contrary to that of Tirole (1996), equilibrium is unique, and lends itself to insightful comparative statics, on the distribution of agents' characteristics as well as on the informational environment.

While the classic game-theoretic approach to (collective) reputation features agents with different behavioral types, Winfree and McCluskey (2005) and Fleckinger (2007) consider oligopoly settings, where group size, or more precisely, the relative weight of an individual agent in the population, plays a central role. Their focus is on the free-riding cost of collective reputation. Fishman et al. (2018) and Neeman et al. (2019) on the contrary highlight some benefits of collective reputation. Both works, rather than generally characterize equilibrium, find conditions under which collective reputation can implement an efficient equilibrium while individual reputation cannot. In Fishman et al. (2018), collective branding of producers of the same type allows consumers to form belief about a brand member's type based on the past performance of all brand members. In a nutshell, the collective channel enables better information through richer sampling. This allows good types to sustain a higher reputation, and even leads to higher prices compared to stand-alone firms.

Neeman et al. (2019) study conditions for a reputation equilibrium under which a competent type always invests. Under collective branding with two firms, consumers do not observe firm-specific but pooled past performance. A collective brand can improve investment incentives in the short-run because its noisier signals prevent reputation from growing too good for incentives. However, in the long run, a collective brand provides less investment incentives than an individual brand. An advantage of our model is that we can conveniently study all equilibria, not only the conditions under which efficient investment can be part of

lies in incentive effects of grouping producers, which differs substantially from the issues analyzed in the literature on statistical discrimination, where groups are given.

[7] It is a 'bad news case', in the wording of Board and Meyer-ter-Vehn (2013). For a detailed analysis of the incentives provided by the informational environment, see Fleckinger et al. (2017).

an equilibrium. Another advantage is that we are not restricted to binary types and one or two agents: This allows to discuss collective reputation design in several dimensions, including the shape of groups and the quality of information, thus paving the way to multiple applications.

In another recent related contribution, Nocke and Strausz (2023) study the moral hazard model inherent to collective brand with multiple joint productions. A global player's contribution complements those of a group of local ones. They study collective branding in the form of pooling information on the output, which induces free-riding on the part of local players. This can however be desirable, as such arrangement is more informative on the effort of the global player. This effect, when sufficiently strong, can dominate local free-riding, which suggests a theory of franchising and optimal brand size. While their model features symmetric local players, heterogeneity of producers is at the core of our approach to group design.

Our comparison of different collective reputation structures also relates to models of group design such as Board (2009). However, while we are concerned with providing incentives for quality, Board (2009) studies how a seller acting as a principal should design consumer groups. The optimal way of dividing players into categories, ex-ante or ex post, in order to provide incentives is also addressed in the literature on status competition (Moldovanu et al., 2007) and optimal grading systems (Dubey and Geanakoplos, 2010). A crucial feature in these works is that the players' payoffs directly depend on their rank or status. In contrast, in our setup producers do not care about status or group assignment *per se*, but only about the informational content provided by the grouping.

Furthermore, our work is related to the literature on Country of Origin (COO). Zhang (2015) explores the strategic dynamics of firms from emerging markets attempting to dissociate from their country-of-origin stereotype, particularly when that COO stereotype implies low quality. When higher-quality firms dissociate themselves from the COO stereotype (by disguising as foreign), they stop contributing to the collective reputation, which can lead to improved overall COO image as the incentives to free-ride of the remaining firms are reduced.

Collective reputation is also analyzed in the substantial literature on brand extension and umbrella branding (e.g. Choi, 1998; Cabral., 2000; Cabral, 2009; Andersson, 2002; Hakenes and Peitz, 2008; Cabral, 2009; Miklós-Thal, 2012). A key difference is that we consider the reputational externalities of pooling heterogeneous producers under a common name, whereas the literature on umbrella branding and brand extension focuses on a firm's profit maximization with coordinated quality and pricing decisions across multiple products, leaving aside the dimension of Bayesian inference at the heart of our model.

Finally, many applications in the literature on collective reputation pertain to agricultural products (e.g. Winfree and McCluskey, 2005; Saak, 2012) or service markets (e.g. Levin and Tadelis, 2005). Recent empirical evidence has also highlighted the effect of collective reputation in times of industry crises (e.g. Bai et al., 2022; Bachmann et al., 2023). The question of comparing the incentive effects of different collective reputation structures has so far not been addressed to our knowledge.

The remainder of the paper is organized as follows. In Section 2, we present our model and determine equilibria under collective reputation as well as under the reference scenario with individual reputations. In Section 3, we compare the quality and welfare brought by collective reputation to what is achievable under individual reputation. Section 4 is dedicated to group stability and the implementation of collective quality with contracts. In Section 5, we discuss admission thresholds in education as an application of our setup. The last section concludes. An extended discussion of the differences between informational environments and the corresponding counterparts to our main results are provided in Appendix A. All proofs are relegated to Appendices B and C.

2 A Model of Collective Reputation for Quality

2.1 Fundamentals

The model features a unit mass of heterogeneous producers, characterized by a marginal cost $\theta \in [0,1]$ of producing quality $q \in [0,1]$. This cost pertains to production technology, individual skills, endowments and so on. It is distributed in the population according to the continuous c.d.f. $F(\theta)$, which is common knowledge. Our theory of collective reputation essentially revolves around how much buyers on the other side of the market know about the characteristic θ of a given producer.

Collective Reputation Structures.

Under our notion of Collective Reputation (CR), each producers belongs to a group, and buyers knowledge is limited to group affiliation and the distribution of costs within that group. Formally, a collective reputation structure allocates producers into a collection of disjoint groups indexed by $i \in \mathcal{I}$. We restrict attention to countable collections of groups $\mathcal{I}$, and for convenience, work directly with the (continuous) distributions over θ within groups.[8] Thus, any countable collection of continuous distributions and associated weights $\{\lambda_i, F_i\}_{i\in\mathcal{I}}$ defines an admissible reputation structure when it satisfies:

$$\sum_{i\in\mathcal{I}} \lambda_i F_i(\theta) = F(\theta) \quad \text{and} \quad \sum_{i\in\mathcal{I}} \lambda_i = 1. \tag{1}$$

This definition is somewhat inspired by information design,[9] since this identity is nothing else than Bayes' rule[10], where λ_i is interpreted as the probability that an agent picked at random belongs to group i. In economic terms, λ_i represents the weight of group i in the population of producers, and F_i is the distribution of types in group i. Importantly, the collective reputation structure is common knowledge, but buyers only observe group membership, and not the type of a producer.

Some widespread forms of collective reputation structures are worth emphasizing. We say that there is Full Collective Reputation (FCR) when there is a single group. Another example, analyzed in Section 5, is the case of an admission threshold, where producers are split in two groups according to a threshold type σ. Furthermore, a reputation structure is horizontal when producers are not grouped according to θ. This is then equivalent to replicating FCR in each group, since the distribution is identically equal to F in all groups. This last case highlights that assuming a continuum of producers effectively neutralizes well-known group size effects.

We will contrast CR with an Individual Reputation (IR) structure, in which buyers perfectly obverse θ. In this case, the set of groups is uncountable, and it is therefore not included in the definition. However, it corresponds to the natural limit where groups become

[8]Our approach can accommodate atoms of producers of some types, but we present the non-atomic case, which corresponds to a smooth density. As regards the countability assumption, this allows us to avoid considerable measure-theoretic technicalities. We however see it as a quite natural assumptions in all the applications discussed.

[9]In doing so, we follow the Bayesian persuasion model of Kamenica and Gentzkow (2011). Forming groups in that way amounts to first add an informational device generating signals that correspond to the group label, according to the mapping of individual producers to groups.

[10]It is hence quite general, despite the countability assumption, which we essentially make for technical ease. It includes for instance, but is not limited to, Borelian structures as in Board (2009), in which all agents of a given type must be in the same group. It includes more general group structures, where two producers with the same costs can belong to different groups. Going beyond countable groups leads to notoriously difficult measure-theoretic issues, similar to the use of continuous information signals in Bayesian persuasion.

arbitrarily small and homogeneous as shown in Section 2.4.[11]

Production and quality signals.

Quality is ex post binary–high or low–and q denotes the probability of obtaining high quality. We sometimes slightly abuse notations by referring to $q = 0$ and $q = 1$ as low and high quality, respectively. A producer with type θ choosing $q \in [0, 1]$ incurs a cost $c(\theta, q) = \theta q$.

On the demand side, buyers are homogeneous, and their willingness to pay is normalized to 1 for high quality and 0 for low quality.[12] The willingness to pay is therefore equal to the expectation of q conditional on the available information. Buyers a priori do not observe the true quality provided by a given producer, but experts provide information on quality through ratings, recommendations, awards etc. The evidence available is positively biased: We assume that high quality is revealed with probability e, which captures the intensity of experts' attention. If quality is low, no news emerge. Hence two kinds of goods are distinguishable: the ones which have been identified as high quality, and the unidentified ones. Finally, a product is traded at its expected quality, i.e. producers are able to charge fully the willingness to pay of the buyers. Importantly, this assumes away competition among producers.

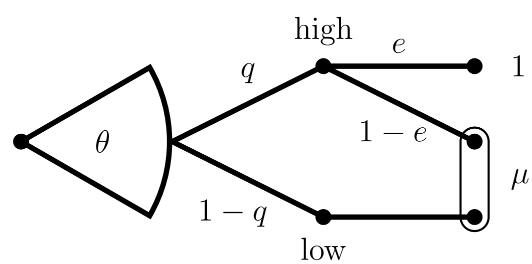

Figure 1: Information structure.

An identified high-quality good is traded at a high price of $p = 1$. Conversely, an unidentified good is traded based on its reputation μ, whether it pertains to the group reputation under CR or the standalone producer's reputation under IR. This reputation for quality μ reflects the belief that the good is of high quality in the absence of a (positive)

[11]Anticipating on the analysis below, we show that the sequence of equilibria under collective reputation where the partition becomes finer and finer, by dividing types into ε length sub-intervals and letting ε go to zero, converges to the individual reputation equilibrium.

[12]Extending the analysis to negative valuations for low quality does not raise difficulties, but can make the set of equilibria richer, see the appendix in Fleckinger et al. (2017).

review from experts. Hence, the term *reputation* pertains only to unidentified products.[13] Unidentified products trade at a price $p(\mu) = \mu.1 + (1-\mu).0 = \mu$.

Equilibrium concept and welfare.

The reputation and quality that emerge result from the incentives created by the interplay of the collective reputation architecture and the informational environment. We adopt the notion of perfect Bayesian equilibrium, where (1) each producer plays his best response to the buyers' belief, and (2) buyers' beliefs are consistent with the actual distribution of quality offered. Since the reputation structure is common knowledge and there is no competition, we can analyze each group in a collective reputation structure separately. For a given group, the buyers' belief, and consequently the incentives of producers within that group, are unaffected by the behavior of producers outside the group. Since we can analyze groups separately, we often refer to the *group equilibrium* for simplicity, when characterizing equilibrium behavior within a given group. Each group contributes to welfare and quality in the perfect Bayesian equilibrium of the full market proportionately to its weight in the CR structure.

Suppose that within a group i all producers with the same type θ choose the same quality $q_i(\theta)$. This is the case in equilibrium, as we will shortly demonstrate, except perhaps for a negligible set of indifferent producers. The aggregate quality is then given by:

$$Q = \sum_{i \in \mathcal{I}} \lambda_i \int_0^1 q_i(\theta) f_i(\theta) d\theta. \tag{2}$$

Aggregate costs C are defined analogously, and the welfare for a given collective reputation structure is:

$$W = \sum_{i \in \mathcal{I}} \lambda_i \int_0^1 (1-\theta) q_i(\theta) f_i(\theta) d\theta. \tag{3}$$

The welfare for the case of IR is defined similarly, by integrating over the individual strategy profile.

Clearly, maximal welfare is attained if all producers choose $q = 1$, since costs are lower

[13] Identified high quality products always have the best reputation, equal to 1. Note that *reputation* refers to expected quality, as in Board and Meyer-ter-Vehn (2013), and not to the type of seller, as in other traditional reputation models. The key distinction lies in the fact that inference is applied to quality, which is endogenous, rather than to a fixed characteristic.

than the social value of quality. This first-best level of welfare is:

$$W_{FB} = \int_0^1 (1-\theta) f(\theta) d\theta = \int_0^1 F(\theta) d\theta. \tag{4}$$

As imperfect information about quality restricts incentives, achieving maximal welfare is typically unattainable.

Incentives for Quality.

For a given reputation μ of a producer, whether their own or that of the group they belong to, depending on the scenario studied, the expected payoff is given by:

$$\Pi(q, \theta) = q(e + (1-e)\mu) + (1-q)\mu - c(\theta, q). \tag{5}$$

The producer can hence always guarantee a payoff of μ, simply by choosing $q = 0$. In turn, for a given reputation μ, a producer with type θ optimally chooses high quality ($q = 1$) if:

$$\theta \leq e(1-\mu). \tag{6}$$

Equation (6) indicates that a good reputation dampens incentives, since the right-hand side decreases with μ. This feature arises as a consequence of the friendly informational environment, and its logic will be applicable in all reputation structures we study below.[14]

2.2 Equilibrium under Individual Reputation

Before analyzing in detail collective reputation, individual reputation constitutes a natural benchmark. Under Individual Reputation (IR), buyers identify each producer individually, so that the cost is public information. Reputation is individual, in the sense that the belief about quality for an unidentified product depends on the identity of the producer, and hence on its anticipated strategy, but not on the others' choices. Facing an individual reputation

[14]The key differences with a hostile environment are presented in Appendix A.

$\mu(\theta)$, a producer chooses $q(\theta)$ to maximize his payoff given by (5), so that:

$$q(\theta) = \begin{cases} 1 & \text{if } \theta < e(1-\mu(\theta)) \\ [0,1] & \text{if } \theta = e(1-\mu(\theta)) \\ 0 & \text{if } \theta > e(1-\mu(\theta)) \end{cases} . \tag{7}$$

Moreover, the reputation $\mu(\theta)$ is consistent when it satisfies Bayes rule:

$$\mu(\theta) = \frac{(1-e)q(\theta)}{(1-e)q(\theta) + (1-q(\theta))} \tag{8}$$

These two conditions together define a perfect Bayesian equilibrium. As seen from (7), incentives are decreasing with reputation, and in particular they vanish if $\mu(\theta) = 1$. It is easy to see that mixing must be involved: producing quality with probability 1 cannot happen in equilibrium if the cost is positive. Indeed, if it were the case, then a buyer would attribute the best reputation to such a producer, but that would in turn destroy incentives. In addition, Equation (7) indicates that if his cost is too high, a producer does not provide high quality. The next proposition characterizes the IR equilibrium.

Proposition 1 (**Equilibrium under Individual Reputation**)
Under Individual Reputation, there exists a unique equilibrium, such that a producer of type θ produces high quality with probability

$$q^*(\theta) = Max\left\{\frac{e-\theta}{e(1-\theta)}, 0\right\}, \tag{9}$$

and the payoff of a producer of type θ is equal to his reputation

$$\mu^*(\theta) = (1-\theta)q^*(\theta). \tag{10}$$

This equilibrium, illustrated on Figure 2, has the expected properties. Higher cost producers choose lower quality, their reputation is lower, and more information increases incentives. In that configuration, welfare is given by $W_{IR} = \int_0^1 q^*(\theta)(1-\theta)f(\theta)d\theta$. Replacing $q^*(\theta)$, then integrating by parts and rearranging yields:

$$W_{IR}(e) = \frac{1}{e}\int_0^e F(\theta)d\theta. \tag{11}$$

The welfare achievable under IR falls strictly below the first-best level when information

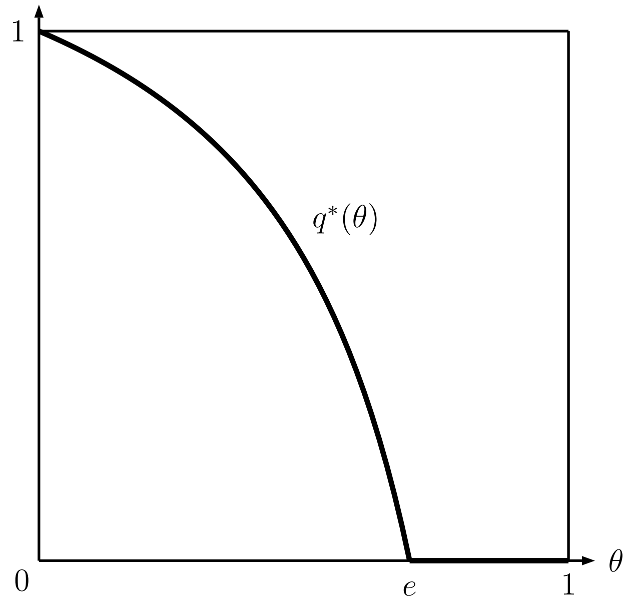

Figure 2: Equilibrium quality under Individual Reputation.

on quality is not perfect. Producers find themselves trapped in a dilemma where their own reputation becomes self-defeating. The temptation to milk a good reputation undermines the creation of powerful incentives. While we consider a simple static framework, mixing in the IR equilibrium mirrors the logic of work-shirk equilibria in good news environments, as in the individual reputation model of Board and Meyer-ter-Vehn (2013), where producers shirk when their reputation is high and work when their reputation is low.

2.3 Equilibrium under Collective Reputation

In contrast to IR, an individual producer alone does not influence the reputation μ_i of his group i, because he belongs to a continuum.[15] Since his payoff is linear in quality, a producer chooses $q = 1$ when $\Pi(1, \theta) \geq \Pi(0, \theta)$, which translates into $\theta \leq e(1 - \mu_i)$ as in Equation (6). An immediate consequence of these incentives is that, in all groups, equilibria have a cutoff structure. They are characterized by cost thresholds θ_i^*, below which producers of group i

[15] By assuming a continuum of producers, our model abstracts from group size effects. This also implies that no individual producer has an influence on the reputation of the group. What matters for the incentive logic at play, however, is not the lack of individual influence on groups, only that the reputation is less responsive to a producer's choice of quality under CR than under IR.

choose high quality (q=1), and above which they choose low quality (q=0):

$$\theta_i^* = e(1 - \mu_i^*). \tag{12}$$

The unidentified goods feature a mix of low and high quality, and the equilibrium reputation obeys Bayes' rule:

$$\mu_i^* = \frac{(1-e)F_i(\theta_i^*)}{(1-e)F_i(\theta_i^*) + (1 - F_i(\theta_i^*))}. \tag{13}$$

Equations (12) and (13) together characterize the equilibrium, as described in the next proposition.

Proposition 2 (**Equilibrium under Collective Reputation**)
For any collective reputation structure $\mathcal{I}$, the equilibrium is unique. In any group $i \in \mathcal{I}$, there exists a threshold θ_i^* such that a producer of type $\theta \leq \theta_i^*$ chooses $q = 1$, and $q = 0$ otherwise. The average quality in group i is

$$F_i(\theta_i^*) = \frac{e - \theta_i^*}{e(1 - \theta_i^*)}, \tag{14}$$

and the reputation is

$$\mu_i^* = (1 - \theta_i^*)F_i(\theta_i^*). \tag{15}$$

In particular, if $F_i(e) = 0$, no producer in group i provides quality, $\theta_i^* = e$, the average quality and the reputation are zero.

Conveniently, the equilibrium is unique, up to possible mixing at the thresholds, which bears no consequence given the continuum of producers. A key observation is that there is *free-riding in all groups* of the Collective Reputation structure, unless information is perfect. This implies that group design alone cannot achieve maximal welfare.

The proposition nests of course the case of Full Collective Reputation (FCR), featuring a single group with weight one and distribution F. The equilibrium characterization is identical within each group, and groups operate independently because competition is absent, so that we can analyze groups separately. In the remainder of the exposition, we therefore drop the index i unless otherwise stated. Put differently, analyzing FCR with a general distribution is equivalent to analyzing any given group.

An illustration of the equilibrium with FCR is provided in Figure 3. The downward-sloping curve depicts the right-hand side of (14), $\frac{e-\theta}{e(1-\theta)}$, which exactly coincides with the

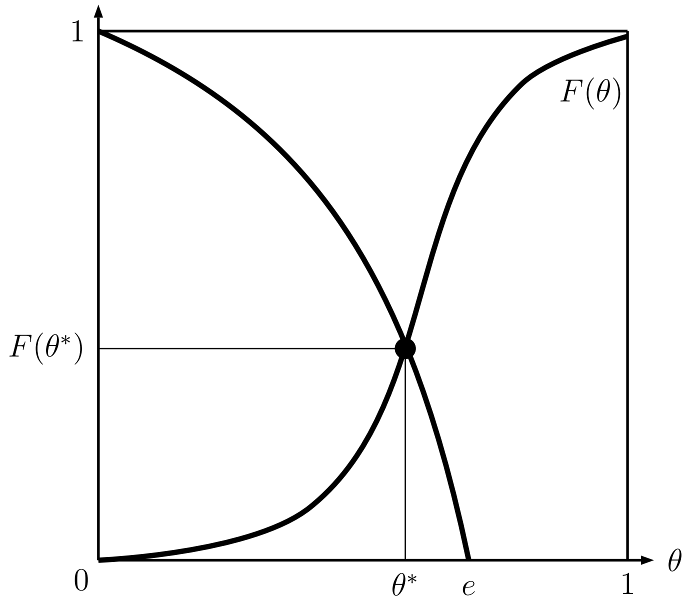

Figure 3: Equilibrium under Collective Reputation.

quality strategy in the the mixed equilibrium under IR, by the indifference property. Indeed, $q^*(\theta)$ corresponds to the average quality produced, such that, when correctly anticipated by buyers, a producer of type θ is indifferent between choosing $q = 1$ and $q = 0$. The equilibrium under FCR is then given at the intersection of this curve with $F(\theta)$, which corresponds to the average quality on the market when all producers with costs lower than θ produce high quality.

The equilibrium welfare under FCR is

$$\begin{aligned} W_{FCR}(e) &= \int_0^{\theta^*} (1-\theta) f(\theta) d\theta \\ &= (1-\theta^*)F(\theta^*) + \int_0^{\theta^*} F(\theta) d\theta \\ &= \mu^* + \int_0^{\theta^*} F(\theta) d\theta. \end{aligned} \tag{16}$$

In particular, the quality of information, measured by e, has both a direct incentive effect and an indirect effect through how it affects collective reputation. In the next section we unpack these two dimensions and undertake a comparison with the case of individual reputation.

Leaning on the analysis in the previous section, we first take a closer look at incentives within groups. As noted above, there is free-riding in all groups of the Collective Reputation structure. Unless information on quality is perfect, a mass of the highest cost producers free-ride on the quality supplied by low-cost producers in the group, and, as a result, reputation cannot reach 1 in equilibrium. How strong is free-riding depends on the composition of the group F, and on the quality of information as captured by e.

The comparative statics regarding the cost distribution determine the incentive effects of group composition. To formalize this, we analyze first-order stochastic shifts in the cost distribution, which offers the appropriate ranking of groups. A cost distribution is considered better than another if its cumulative distribution exceeds the other's at every point in the cost range. The comparative statics regarding information on quality is subsumed in the parameter e, and represents for instance changes in the intensity of experts' attention. In the following proposition, we omit the group index i since, as explained above, the analysis is identical across groups.

Proposition 3 (Comparative Statics with Collective Reputation)
Consider ***a stochastic reduction in costs*** *(i.e. a first-order stochastic decrease in F).*
(1) Average quality increases: $F(\theta^*)$ increases, (2) Reputation increases: μ^* increases and (3) Incentives in equilibrium decrease: θ^* decreases.

Consider ***an increase in information on quality*** *(i.e. an increase in e).*
(4) Incentives and average quality increase: θ^*, and consequently $F(\theta^*)$, increase. More information has ambiguous effects on reputation: (5a) μ^* increases for e close to 0, and decreases for e close to 1 and (5b) if, in addition, F is log-concave, μ^* is single-peaked in e: reputation first increases then decreases with e.

Figure 4 illustrates how the group equilibrium is affected when costs are reduced and information is improved. In particular, more producers provide quality in a better group, although incentives to free-ride become stronger. Conversely, improved information enhances both the quality provided and incentives simultaneously.

A direct consequence of the proposition is that every producer would like to expel the (other) high-cost producers who free-ride, to obtain a higher reputation μ^*. However, expelling high-cost producers lowers incentives in the group in return, by increasing its reputation. Conversely, if low-cost producers (below θ^*) are separated from the group, incentives of the others increase.

Proposition 3 also indicates that all types of producers, regardless of whether they provide quality or not, value *some* information due to its positive impact on reputation. Those who

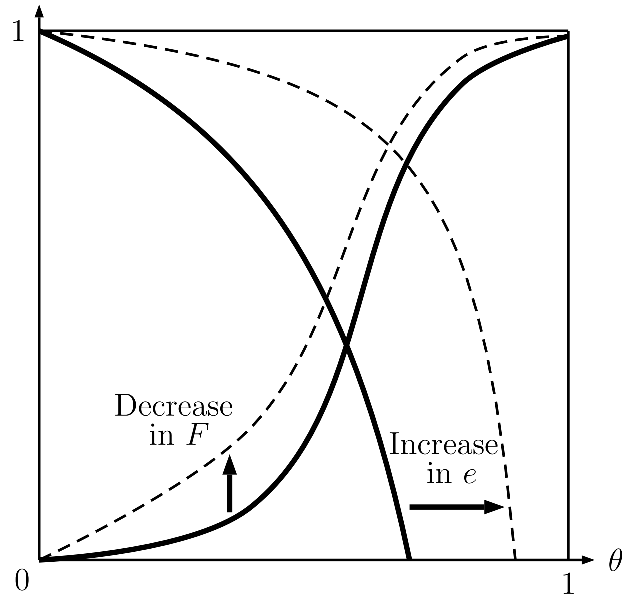

Figure 4: Comparative statics of the Collective Reputation equilibrium.

provide quality obtain a payoff $e+(1-e)\mu^*-\theta$: they always benefit from better information because it both enhances their chances of recognition and boosts the group's reputation. Conversely, those who free-ride can benefit because they obtain a payoff equal to μ^*, which initially increases with e. Notably, when F is log-concave, producers unanimously agree to information improvements until the reputation reaches its peak.

2.4 Individual Reputation as a limit of Collective Reputation

Our central task in the next section will be to compare IR to CR. To make this comparison meaningful, we first point out that IR is, reassuringly, compatible with our notion of collective reputation. The next proposition shows indeed that IR can be attained as the limit of CR structures where producers are grouped into finer and finer cost intervals. It hence confirms natural intuition: by grouping more and more homogeneous producers, the limit CR structure approaches IR, in which θ is know to buyers, and the corresponding equilibrium quality and welfare do converge to the IR levels.

Proposition 4

The Individual Reputation equilibrium obtains in the limit of a sequence of Collective Reputation structures. Quality, costs and welfare in the corresponding sequence of equilibria uniformly converge to the IR equilibrium values.

The IR equilibrium belongs to the closure of the set of equilibria resulting from CR structures. Despite the markedly different equilibrium behavior under IR and CR, a well chosen CR sequence leads to the same aggregate behavior in the limit. Seen from the other direction, this exercise can be regarded as purification of the mixed equilibrium under IR. Noticeably, the convergence is uniform in the sense that, in each group of the CR structure converging to perfect knowledge of the cost, quality and costs converge to that obtained under IR over the corresponding interval of costs.

3 Collective vs. Individual Quality

Collective reputation has countervailing incentive effects, as demonstrated by the comparative exercise in Proposition 3. In particular, excluding bad types (high-cost members) increases reputation, but reduces incentives. Conversely, excluding good types (low-cost members) reduces reputation, but increases incentives. To highlight the efficiency effects of collective reputation and trade-offs inherent in different collective reputation structures, we compare the two extreme structures, Full Collective Reputation and Individual Reputation. As noted after Proposition 2, the restriction to FCR is harmless, since the problem in a group simply correspond to FCR with a particular cost distribution, and groups are independent. Comparing FCR to IR is equivalent to asking when transparency on costs for a given group is desirable. As we shall see, the central trade-off is between more low-cost producers providing maximal quality under FCR, versus more quality from a larger pool of producers under IR.

Under FCR, the group's reputation induces high-cost producers to free-ride. However, under FCR, low-cost producers ($\theta \leq \theta^*$) all provide the highest quality. In contrast, under IR, a producer with a positive cost never provides maximal quality, because the temptation of reputation-milking partially destroys incentives. In particular, under IR, all producers with $0 < \theta < e$ choose quality $0 < q < 1$. Thus, the more efficient types ($\theta < \theta^*$) exert more effort under CR than under IR. However, under CR producers of type $\theta > \theta^*$ choose the lowest quality, while under IR types θ with $\theta^* < \theta < e$ choose $q > 0$: there is more effort under IR than under FCR from the less efficient types. Put together, effort allocation is thus less efficient under IR than under CR. This implies that *for the same given level of quality*, costs under IR are higher than under FCR. An immediate consequence is that IR can generate higher welfare than CR only if it generates higher average quality. The next proposition delineates the efficiency benefits of CR.

Proposition 5 (Full Collective Reputation vs Individual Reputation)

1. *For a given positive total quality, FCR achieves it at strictly lower total cost than IR.*

2. *If total cost is higher under FCR than under IR, then quality is higher under FCR.*

3. *If quality is higher under FCR than under IR, then welfare is higher under FCR.*

A significant advantage of collective reputation over individual reputation arises from a better allocation of quality effort: FCR yields a cost efficiency benefit. Additionally, if FCR generates higher quality, it necessarily surpasses IR, due to the combination of higher quality and cost efficiency. However, IR can generate a higher level of quality, in which case the comparison becomes ambiguous. As is evident from the characterization of equilibrium under collective and individual reputations, the comparison of quality depends finely on the properties of the cost distribution.

Consider a reduction in costs, i.e. a first-order stochastic decrease of the cost distribution. Under IR, producers' incentives are unchanged, and the only effect is positive: the relative weight of more efficient producers providing some quality increases. In turn, under FCR, such a reduction in costs increases in addition the reputation of the group, which reduces incentives for all producers. While this might suggest that a reduction in costs makes IR relatively more desirable than FCR, this is not necessarily true. Indeed, as low costs producers provide full quality under FCR, their increasing weight with a first-order stochastic decrease can more than compensate the incentives reduction due to a higher reputation. Overall, the gains with FCR can still be higher than under IR.

An interesting question is how collective reputation fares compared to individual reputation depending on the level of market information e. The main trade-off when increasing the quality of information is that the gains are spread across the population under IR, while they are concentrated around the equilibrium threshold under CR. Increasing information increases welfare under both arrangements, but whether one or the other arrangement leverages better the information improvement depends on the comparison of additional incentives in the whole population with IR, versus an increasing set of agents providing quality under CR.

As it turns out, the comparison is clear-cut when we look at the limit cases where information is either very poor or almost perfect, under the following mild assumption on the distribution.

Definition 1 (**Regular distribution**)
A distribution F is regular if its density f is continuous, positive and bounded.

The next proposition shows that when e is sufficiently small, FCR does strictly better than IR in terms of welfare for any regular distribution.

Proposition 6 (**FCR vs IR with poor information**)
In the limit $e \to 0$, welfare is weakly higher under FCR than IR :

$$W_{FCR}(0) - W_{IR}(0) = 0 \quad \text{and} \quad \lim_{e \to 0} \frac{d}{de}\left[W_{FCR}(e) - W_{IR}(e)\right] \geq \frac{f(0)}{2}.$$

If F is regular, Full Collective Reputation hence strictly dominates Individual Reputation if e is small enough.

When information about quality is very poor, collective reputation becomes essential. Neither FCR nor IR generate any quality when $e = 0$, but introducing even a slight chance of good news–for instance, a small gesture of goodwill from experts–triggers a stronger incentive response under FCR.

This result is noteworthy: When e is low, the market's asymmetric information problem is particularly severe, and one might initially expect that the added free-rider problems would undermine collective reputation structures. However, the proposition shows the opposite: free-riding across producers under collective reputation actually helps. Intuitively, the effect of reputation is strong when e is low, since its weight in a producer's payoff that provides quality is $(1 - e)$. Under FCR, the reputation remains low for everyone, as most producers do not provide quality, while under IR, the most efficient producers' milking temptation is not mitigated by a comparable effect.

At the other extreme, the next proposition studies the case of almost perfect information.

Proposition 7 (**FCR vs IR with almost perfect information**)
Suppose that F is regular. Welfare under FCR is strictly higher than welfare IR if e is high enough and the average cost is above 1/2. Formally,

$$W_{FCR}(1) - W_{IR}(1) = 0 \quad \text{and} \quad \lim_{e \to 1} \frac{d}{de}\left[W_{FCR}(e) - W_{IR}(e)\right] = \frac{1}{2} - \mathbb{E}(\theta).$$

The comparison between FCR and IR is ambiguous when information is almost perfect. FCR dominates if the cost is high enough, while IR dominates otherwise. Of course, both arrangements yield the first-best outcome when $e = 1$. However, they reach this limit differently as e goes to 1. The welfare under IR increases proportionately to average costs, which reflects that better information increases incentives for all producers. Conversely, under FCR, only marginal producers at the threshold matter, in this case those with costs close to 1. Even though they produce very low welfare given their high costs, they are likely to weight more the higher the average costs. The trade-off is reflected surprisingly sharply in the proposition.

Uniform distribution.

Whether collective reputation yields higher welfare than individual reputation depends both finely on the type distribution of producers as well as the level of market information. This is not surprising as incentive effects operate at the 'extensive margin' in one case (around the threshold, with FCR), and at the 'intensive margin' in the other case (over the whole population of quality providing producers, with IR). To better illustrate the incentive and efficiency implications of collective versus individual reputation, we will delve into the case of a uniform costs distribution. With this distribution, closed-form solutions are readily available and derived in Appendix C. Figure 5 shows equilibrium qualities in both structures, and highlights the quality and cost trade-off between both structures. Quality under FCR is given by the area of the shaded rectangle, and quality under IR by the dotted area. The next proposition shows that FCR always dominates IR in terms of welfare when the cost distribution is uniform.

Proposition 8 (**FCR vs IR with the uniform distribution**)
Suppose that F is the uniform distribution. Then Full Collective Reputation yields higher welfare than Individual Reputation.

In the case of the uniform distribution, Full Collective Reputation attains higher welfare than Individual Reputation for any level of information. In the example shown in Figure 5 with $e = 2/3$, visual inspection indicates that quality is higher under FCR than under IR. When this is the case, Proposition 5 already implied that FCR produces higher welfare than IR. Proposition 8 asserts in addition that for the uniform distribution, welfare is always higher under FCR, even for levels of information for which IR generates higher quality.

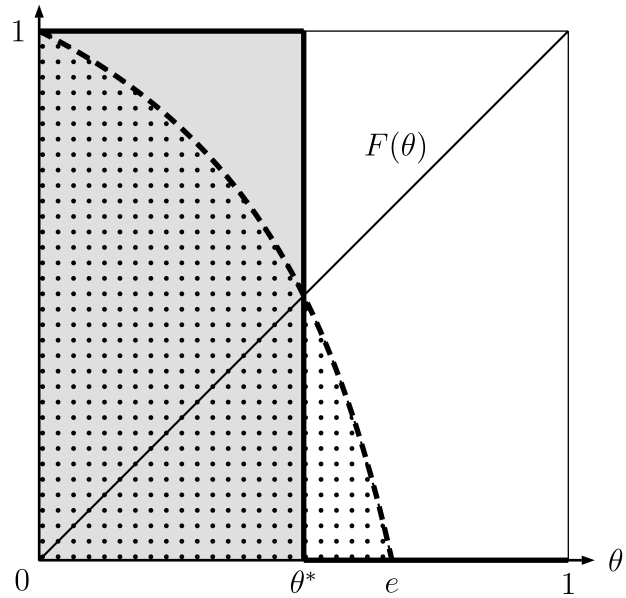

Figure 5: Quality provision (uniform distribution). Solid lines and grey area pertain to FCR, dashed lines and dotted area to IR.

Figure 6 displays quality, costs and welfare of FCR and IR as functions of e, with the uniform distribution. Depending on e, several configurations emerge. If e is high enough, quality is higher under IR, but the higher costs still makes FCR preferable to IR in terms of welfare. It can even be the case simultaneously that quality is higher and costs lower under FCR, for intermediate information levels.[16]

4 Sustaining Collective Quality

4.1 Collective Reputation Unraveling

As shown in the previous section, welfare and payoffs for a producer organization that features some form of collective reputation—or even full collective reputation—can be higher than under individual reputations. However, this does not imply that, individually, producers would prefer to stay in a group rather than work under their individual reputations. In fact, individual incentives lead the lowest-cost producers in any quality-producing group to prefer

[16]More precisely, using the calculations in Appendix C, one obtains that when $e \in [.7309, .8335]$, aggregate quality is higher and total cost is lower with FCR than with IR. The uniform case hence illustrates that even the most counter-intuitive configuration allowed by Proposition 5 can arise in non-pathological circumstances.

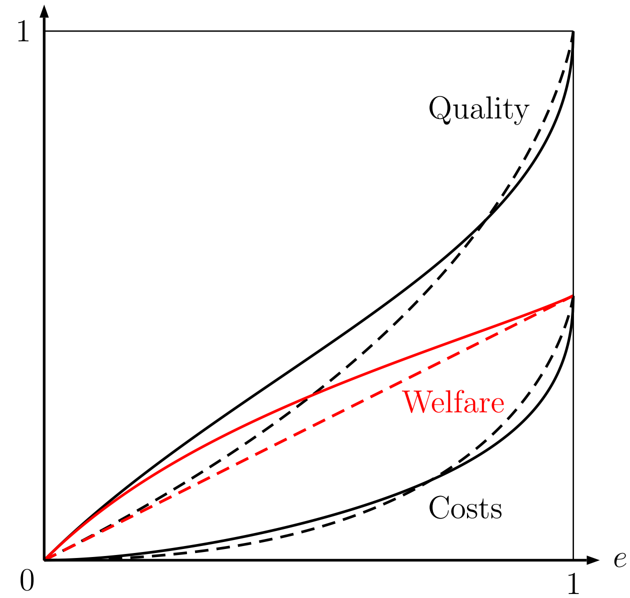

Figure 6: Quality, Costs and Welfare (uniform distribution). Solid lines pertain to FCR, dashed lines to IR.

breaking away, causing the collective reputation to inevitably unravel. This is shown in the next proposition.

Proposition 9 (**Group unraveling**)
In a collective reputation structure, any group producing quality unravels to individual reputation.

Unraveling occurs because, in any group, the lowest cost producers always obtain a higher payoff under individual reputation. This payoff difference is illustrated in Figure 7 for FCR with the uniform distribution. Note that, contrary to unraveling with pure adverse selection, which is efficient, unraveling can be inefficient because of moral hazard. Figure 7 precisely displays a case in which FCR yields higher welfare (and producer surplus) than IR. Indeed, a direct corollary to Proposition 8 is that unraveling with the uniform distribution is suboptimal.

This raises the question of whether a group can implement an internal incentive system to sustain collective reputation in order to reap its benefits. Clearly, when costs are observed within a group and transfers between group members can be individualized, the benefit of collective reputation can be redistributed to low-cost producers to entice them to stay in the group. What we will show in the next section is that even when costs are not observed within the group and transfers cannot depend on costs, collective reputation can still be sustained.

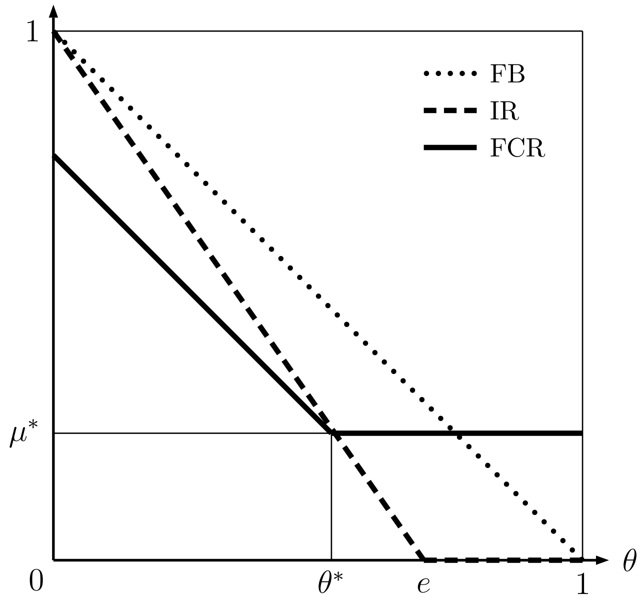

Figure 7: Individual payoff under FB, IR and FCR.

More strikingly, we show that a simple uniform transfer scheme can ensure participation and increase incentives of group members to the First Best quality level.

4.2 Efficient Group Contracts

We consider group contracts signed with a Collective Organization (CO) bilaterally by each member. The only information available for contracting is the market price received individually by producers. When types are not observed in a group, one can easily see that the only mechanisms available must be pooling, i.e. it is impossible to tailor incentives to the type of a producer when it is not observable. A contract comprises (at most) two instruments, since there are only two possible market prices.

We will consider, without loss of generality, the following group incentive scheme comprising two instruments. Initially, a participation fee t is levied on each producer. After market clearing, producers with identified high quality receive a bonus b, paid by the CO on top of a high market price. Equilibrium quality depends on the bonus s, but t does not influence the equilibrium, once participation is ensured. Let θ_b and μ_b denote the equilibrium threshold and reputation with a bonus b. The incentive constraint now writes:

$$e(1+b) + (1-e)\mu_b - \theta - t \geq \mu_b - t. \tag{17}$$

In addition, Bayes rule determines the reputation for the corresponding threshold as in (13). Combining these two relations yields the following new equilibrium relation taking b into account:

$$F(\theta_b^*) = \frac{e(1+b) - \theta_b^*}{e(1+eb-\theta_b^*)} \tag{18}$$

In order to ensure participation of a producer of type θ, the scheme should provide a higher payoff than if the producer decides not to join, i.e., higher than the IR payoff:

$$e(1+b) + (1-e)\mu_b^* - \theta - t \geq \mu^*(\theta). \tag{19}$$

Finally, the scheme is (collectively) budget-balanced if the participation fees levied cover the bonuses paid:

$$t \geq eb. \tag{20}$$

Stability and budget-balance being defined, we can now state the next result.

Proposition 10 (**Bonus club**)
Suppose the producers each have at least one unit of capital. The following scheme is stable, budget-balanced and implements the first-best: the producers pay a fee $t^* = 1$ to join the group, and receive a bonus subsidy $b^* = \frac{1}{e}$ for an identified high quality product.

In a bonus club, the CO sells incentive lottery tickets to producers. The ticket price t entitles them to sell under the collective reputation, and to collect the bonus b with the ticket if their quality is recognized externally. Concretely, producers market the product themselves under the group name, and an explicit reward for quality on top of market price is attached to the (costly) membership to the collective organization. When adhering to this bonus club scheme, producers all receive in expectation their first-best payoff of $1-\theta$, which they cannot attain individually.

Importantly, the bonus club scheme requires capital from the producers ex-ante, on top of the sunk cost θ. An alternative payoff-equivalent scheme does not require membership fees. It consists in producers delivering their products, instead of capital, to a CO that is also in charge of marketing all the products, such as a cooperative. The CO can then redistribute with distorted prices, to calibrate incentives. The CO can indeed not only keep the proceeds from the basic reputation price μ to increase the incentive gap, but it can in addition redistribute to high quality producers to increase incentives even further.

Proposition 11 (**Collective retail channel**)

The following scheme is stable, budget-balanced and implements the first-best: the producers supply a collective retailer, that markets the products and pays back the producers 0 *for unidentified quality and* $\frac{1}{e}$ *for good quality.*

Proof. Participating in the group amounts here to obtaining a lottery ticket that is worthless without effort, but worth 0 with probability $(1-e)$ and $1/e$ with probability e when providing quality. In total, participation with effort hence yields the first-best payoff $1 - \theta$, which beats the alternative of not participating, and efficiency is maximized. The scheme is clearly budget-balanced since all prices on the market are 1 in equilibrium, which covers the total amount redistributed. ■

The main advantage of a collective retail channel, compared to a bonus club, is that it does not require ex-ante financing. Both have otherwise very comparable properties, in particular that of screening out potential producers with costs higher than 1, who would gain nothing from participating.

One concern with a collective retail channel is that it has ex-post incentive to pretend that quality was not recognized on the market. The bonus club scheme is on the contrary susceptible to manipulation by producers, that have an additional interest in pretending their quality was recognized as high to collect the bonus. Overall, though, a bonus club requires only a minimal CO, that simply acts as a bonus bank, while a collective retail channel empowers a new logistics and sales intermediary.

To conclude this section, we would like to point out that both proposed schemes aim at increasing incentives by redistributing the benefits of collective reputation. This tends to increase risk and, in doing so, also amplifies inequalities ex-post, even though all types of producers are treated equally ex-ante. The bonus club scheme illustrates starkly that only the producers already benefiting from high prices also benefit from the advantages of the collective organization. Interestingly, as shown by Gergaud et al. (2017), this turns out to be a feature of the collective organization for Bordeaux wines–admittedly a collectively successful region. Likewise, famous alumni are more often praised and rewarded by their alma mater than average students.

5 An Application to Admission Thresholds in Education

Education institutions create distinct groups as they admit students and award degrees. Students seek to join highly reputed schools, and the reputation of a school depends in turn on the selectivity of entry in the pool of students it admits. What type of incentives does this kind of grouping generate? How should a public education system, with a welfare criterion in mind, organize to leverage the benefits of collective quality? What types of business models do private institutions devise to extract surplus from their admission standards? Our model is well-suited to explore these questions. While human capital acquisition and peer-effects among students do play an important role in reality, we use our framework to focus on the reputation stemming from group composition. Ignoring other externalities parallels assuming away competition among producers in the main interpretation developed so far. Our approach to education systems overall depicts a purely informational scenario, where the strategic interaction between students only goes through their group reputation.

Collective reputation structures with an admission threshold.

Schools administer tests to allocate a limited number of seats to an heterogeneous set of applicants. A straightforward approach to modeling an admission threshold is to divide the population of students (producers) into two groups based on their cost parameter. Let σ be the admission threshold, that is, the maximal cost allowing admittance in the top group.[17] We assume in addition that students cannot operate under individual reputation, i.e. they cannot disclose their type individually and work under IR. Note finally that with two groups, it does not matter whether group 2 formally exists or not: students of group 2 are simply those not in the formal group 1.

Group 1 gathers all students with $\theta < \sigma$: as a convention, students with $\theta = \sigma$ belong to group 2.[18] An admission threshold generates a CR structure $\mathcal{A}(\sigma)$, containing two groups,

[17] This implicitly assumes that the entry test is perfect. Our definition of collective reputation structures allows to entertain imperfect admission threshold, because the set of groups does not need to be Borelian: students with identical types can belong to different groups. An imperfect test still amounts to splitting the population according to a pass/fail entrance test, which generates two groups ordered by first-order stochastic dominance. The comparative statics on groups, and hence the mechanism studied below, are qualitatively the same as with a perfect admission threshold, thanks to Proposition 3.

[18] The highest cost students in a group never exert effort, and the continuum makes it a harmless convention. In addition, both extreme cases, i.e. $\sigma = 0$ or $\sigma = 1$, collapse to FCR.

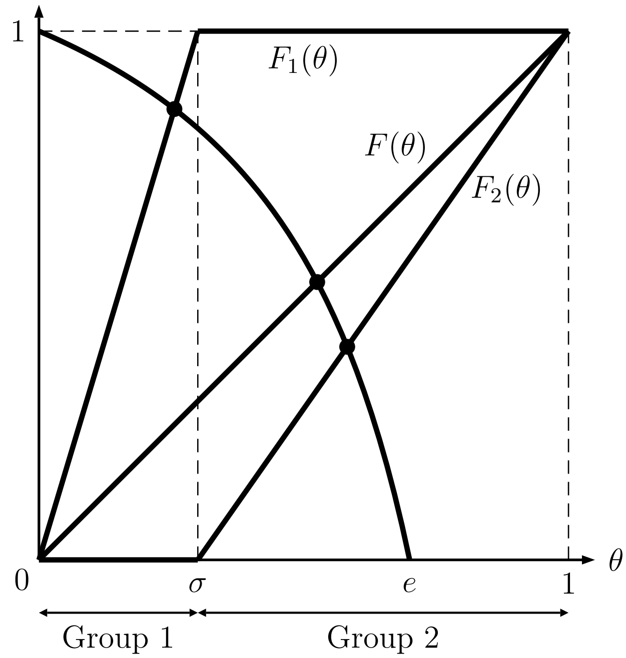

Figure 8: Equilibrium with an admission threshold (uniform distribution).

with truncated distributions defined by:

$$F_1(\theta) = \min\{\frac{F(\theta)}{F(\sigma)}, 1\} \text{ and } F_2(\theta) = \max\{\frac{F(\theta) - F(\sigma)}{1 - F(\sigma)}, 0\}, \tag{21}$$

and associated weights

$$\lambda_1 = F(\sigma) \text{ and } \lambda_2 = 1 - F(\sigma). \tag{22}$$

To obtain the equilibrium under this CR structure, Proposition 2 applies. The result is represented on Figure 8.

The groups' distributions have a number of properties worth keeping in mind. For any σ, F_2 first-order stochastically dominates F, which first-order stochastically dominates F_1. For any cutoffs σ and σ' with $0 < \sigma < \sigma' < 1$ and associated distributions F_i and F_i', F_i' first-order stochastically dominates F_i. That is, an increase in σ stochastically increases costs in *both* groups.

Combined with the equilibrium comparative statics obtained in Proposition 3, these properties have several implications. Average quality in equilibrium is ordered intuitively: $F_1(\theta_1^*) > F(\theta^*) > F_2(\theta_2^*)$, but incentives are lower in the first group: $\theta_1^* < \theta^* < \theta_2^*$. Since reputation increases for group 1 students compared to the case of a single group, free-riding is comparatively more tempting, by Equation (12). Compared to FCR, $\mathcal{A}(\sigma)$ decreases incentives for students in group 1, but it creates some for a group of less talented students, now in group 2. This also implies that the allocation of effort among students is not optimal: those with costs between θ_1^* and σ do not produce quality, while less efficient ones, with costs between σ and θ_2^*, do produce quality. Whether an admission threshold increases welfare compared to FCR depends on whether the incentive effects in group 2 can overcome both the reduced incentives in group 1 and the cost misallocation.

Efficient admission threshold.

We slightly abuse notation and define the welfare in equilibrium under $\mathcal{A}(\sigma)$ as:

$$W(\sigma) = \lambda_1(\sigma) \int_0^{\theta_1^*} (1-\theta) f_1(\theta) d\theta + \lambda_2(\sigma) \int_0^{\theta_2^*} (1-\theta) f_2(\theta) d\theta \tag{23}$$

Perhaps surprisingly, the next result shows that it is always possible to pick an admission threshold that strictly increases welfare, compared to Full Collective Reputation.

Proposition 12 (**Admission Threshold**)
Suppose that F is regular. There exists an admission threshold $0 < \sigma^* < e$ such that $W(\sigma^*) > W_{FCR}$.

With a regular[19] distribution, FCR can be improved on by introducing selective admission into an elite group. Note that, in contrast, choosing a cutoff σ with $e \leq \sigma < 1$ yields a strictly lower welfare than FCR. Indeed, for $\sigma \geq e$, group 2 generates no quality, and, since $\theta_1^* < \theta^*$, quality is lower than under FCR. As a result, the welfare under $\mathcal{A}(\sigma)$ is also lower when $e \leq \sigma < 1$.

The welfare increases with σ when $\sigma \geq e$, since θ_1^* increases with σ, while quality is always absent in group 2. Adding students providing no quality to group 1 decreases reputation, which in turn increases incentives at the top. Put differently, an admission threshold can backfire if it is chosen too loose: in this case, no selection at all is a better option. Figure 9 illustrates this in the case of the uniform distribution.

The general question of splitting or merging groups bears an apparent resemblance to the classic issue of whether third degree price discrimination increases welfare, dating back

[19] In the proof of Proposition 12, we only use the fact that $f(0) > 0$ for a regular distribution. We discuss this point below.

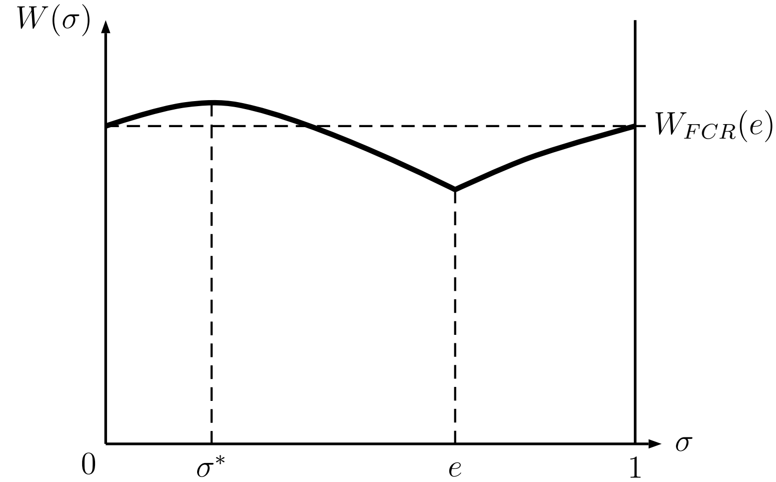

Figure 9: Welfare with an admission threshold (uniform distribution).

to Robinson (1933), and extensively studied since then (see Schmalensee, 1981; Varian, 1985, for important milestones). However, the mechanism is quite different, although the type distribution plays a central role in both cases, whether that of students or producers costs here, or that of consumers' taste in a third degree price discrimination model. Our strategic framework is governed by a different logic, featuring a Bayesian reputation channel and incentives concerns, and the insights and techniques from models of third degree price discrimination do not carry through to our setting.

To conclude the analysis of optimal admission thresholds, Proposition 12 implies that a strong form of elitism is desirable with a regular distribution. Indeed, the result can be iterated *ad infinitum* for any group containing zero-costs students. Since total welfare is a weighted sum of welfare in each group, improving repeatedly the first term in this convex combination is desirable. For instance, starting with FCR, a series of well-chosen admission thresholds improves welfare, by repeatedly applying Proposition 12 to the top group. This might seem puzzling at first sight, since FCR can be strictly better than IR, as we have seen for instance in Proposition 8 with the uniform distribution. This is not a contradiction, as the iteration of the result above applies only at the top with zero costs, and not over the whole cost interval.

This welfare-improving iterated use of admission thresholds also shows that no finite CR structure can be welfare-optimal. What would be an optimal CR structure is a topic left for further research. Instead, we study below a more concrete application of our model to a private education monopoly.

Profit-driven selection.

An admission threshold can improve welfare, but the selectivity of admission depends on the motive for creating groups. We compare now selection by a private education provider to the efficient selection level. Proposition 12 describes the socially optimal university structure, $\mathcal{A}(\sigma^*)$, where σ^* maximizes welfare in a two-tier system. For instance, group 1 and 2 correspond, respectively, to students with or without a college education. Would a private university with market power choose the same admission threshold, and if not, what kind of distortion can be expected? To answer these questions, we consider a profit-maximizing school, that chooses an admission threshold σ and a tuition fee t.

As a preliminary step, we need to define the gain that a given student with type $\theta < \sigma$ obtains when paying the tuition fee t to join the private school.[20] We assume that the alternative, joining group 2, is free, which is consistent with either no college education, or fully subsidized college education. For the structure with two groups studied above to result in an equilibrium as described in Proposition 12, it must be the case that all students in group 1 prefer not to join group 2. In other words, all students that are admitted to group 1 must be ready to pay the corresponding tuition rather than remain uneducated (or educated with high-cost students). Since $\sigma < \theta_2^*$, all students that can be admitted to group 1 would choose $q = 1$ in group 2. Hence, for the best students, those who choose $q = 1$ when they belong to group 1, to join the private university, it must be the case that $e + (1-e)\mu_1 - \theta - t \geq e + (1-e)\mu_2 - \theta$, or $t \leq (1-e)(\mu_1 - \mu_2)$. For free-riding students in group 1, i.e. with type $\theta_1^* < \theta < \sigma$, this requires $\mu_1 - t \geq e + (1-e)\mu_2 - \theta$, or $t \leq (1-e)(\mu_1 - \mu_2) + \theta - e(1-\mu_1)$. The additional term is positive because the incentive constraint of these students is not met in group 1. The participation constraint of students with $\theta \leq \theta_1^*$ is therefore tighter.[21] As a result, for any admission threshold σ, the private university optimally chooses a tuition level:

$$t(\sigma) = (1-e)(\mu_1(\sigma) - \mu_2(\sigma)), \tag{24}$$

which implements the equilibrium where students with $\theta < \sigma$ are willing to join the private

[20] We keep the harmless convention that types $\theta = \sigma$ cannot be admitted to the private school.

[21] Notice that a private university is tempted to discriminate students with flexible tuition, depending on their θ. We ignore here this aspect, whereas in reality it is common practice to grant fellowships to top students, thereby reducing their tuition, while at the same time admitting legacy students, who pay higher tuition but may not have low enough costs to be incentivized under the high reputation of group 1.

university. This results in a profit equal to the total tuition proceeds from group 1 students:

$$T(\sigma) = (1-e)(\mu_1(\sigma) - \mu_2(\sigma))F(\sigma). \tag{25}$$

This expression deserves a couple a observations. First, private universities can only be profitable because information on quality produced is imperfect. In the limit to perfect information, with e close to 1, most students provide quality, and it is detected so often that collective reputation barely plays a role. Hence, the gains in joining the private university are low, and so must be tuition fees and profits in this case. Secondly, the monopoly problem consists of selling a quantity of seats $F(\sigma)$ with reputation increments $(\mu_1(\sigma) - \mu_2(\sigma))$, together with attracting the best students, who provide quality and are building the reputation of the school. Applying Proposition 3 to the truncated distribution indicates that both $\mu_1(\sigma)$ and $\mu_2(\sigma)$ decrease with σ. However, $\mu_1(\sigma) - \mu_2(\sigma)$ has no reason to be monotonic, and the monopoly problem is more intricate than may seem at first.

As regards the effect of the information, Proposition 3 asserts for instance that if F is log-concave, μ_i is single-peaked in e. This implies that even with a log-concave F, the university profit is unlikely to be monotonic in e for a given admission threshold. The effect of the level of information on the tuition-maximizing threshold is, as a result, nontrivial.

While finding a tuition-maximizing threshold $\sigma_T(e)$ is an intricate issue in general, we present numerical results with the uniform distribution,[22] as represented in Figure 10, along with the efficient admission threshold, $\sigma^*(e)$.

The tuition-maximizing threshold deserves some explanations, as it follows four regimes, delimited by the vertical thin lines. For low e's, a tuition-maximizing school fully privatizes the market, with $\sigma_T = 1$. It becomes selective when e is higher, in the second regime, but the corresponding threshold is such that $e < \sigma_T < 1$, which is inefficient, as discussed after Proposition 12. In the third regime, the private school even chooses the worst possible threshold from a welfare point of view: $\sigma_T = e$. In this case, the private institution makes sure that the reputation of group 2 remains equal to zero. Increasing the threshold from e would bring more students, but would reduce the reputation, and hence the tuition. Conversely, being more selective allows to raise the tuition, but also allows reputation of group 2 to develop, which backfires on the relative attractivity of group 1 and reduces the feasible

[22] The limit cases with $e = 0$ and $e = 1$ are excluded. In both cases, any threshold yields the same level of tuitions and welfare. All computations are made based on the last section of the Appendix, pertaining to the uniform distribution.

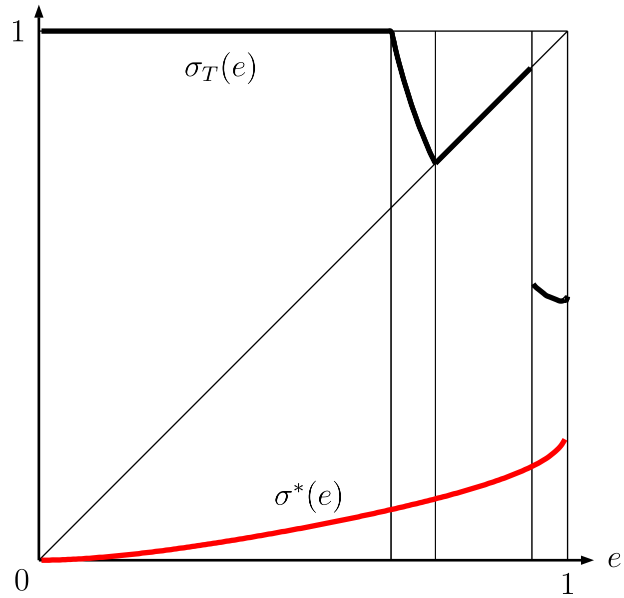

Figure 10: Efficient vs tuition-maximizing thresholds (uniform distribution).

tuition. Finally, for e close to 1, in the fourth and last regime, it is very costly to keep group 2 reputation to zero, and the private institution prefers as a result to be selective and increase tuition. The optimal threshold is non-monotonic in this region: it first decreases then increases, to approach $\frac{1}{2}$ from below in the limit to perfect information. When $e = 1$, collective reputation is irrelevant and the first-best is attained for any threshold, while tuitions are zero, as reputation plays no role.

We have abstracted from an essential part of the education system: it supposedly builds human capital. The model features neither teaching and infrastructure costs for the university, nor educational benefits for students, except the ones created by collective reputation. Of course, a full-fledge model of education should incorporate them. Our bare-bone model revolves around the signaling value of education in the vein of the seminal contribution by Spence (1973), with the added twist of quality production by students in the education system. To that extent, our model probably suits best masters and PhD granting institutions, which students productions have intrinsic value.

6 Conclusion

Collective reputations are pervasive in many industries. We develop a simple model of collective reputation for quality and show that collective incentives are well-suited to address informational problems that plague the exchange of quality products. Our analysis highlights two key advantages of collective reputations over individual ones: First, it reshapes incentives in a way that can lead to more cost-efficient production. In particular, we emphasize the incentive benefits of some free-riding, and show that it may be less harmful to efficiency than reputation-milking in systems based on individual reputations. Second, it enables the use of group-based incentives on top of market incentives–without any additional information requirements–to enhance quality provision.

We demonstrate that even without explicit group incentives, pooling reputations can be an effective mechanism for ensuring quality. This approach helps mitigate the issue of quality under-provision caused by imperfect monitoring and, in some cases, can achieve optimal outcomes when transfers are feasible. Notably, we find that collective reputation is particularly advantageous in situations where quality monitoring is most challenging.

The application of our model to admission thresholds in education systems underscores its flexibility and potential for further extensions, particularly in the design and merging of groups. Replicating dynamic reputation effects in a tractable static model, our framework can be adapted to a wide range of contexts, from the coexistence of public and private tiered systems in higher education to the splitting of wine producers from traditional appellations in France and Italy. Since a well-designed collective reputation structure can generate higher welfare than individual reputation systems, intermediaries such as retail platform and online services can also be analyzed through the lens of our model. These intermediaries, with their new business models, have reshaped the economy by capturing the value created by collective reputation. Our framework opens new avenues to understand their strategies in terms of the joint design of rating systems and product grouping.

An important aspect that we have kept exogenous is the intensity of attention paid by experts. If we allow this attention to depend on the composition of groups, additional interesting phenomena arise. Suppose, for instance, that experts have a slight preference for discovering high quality—perhaps because it enhances their audience and revenue, or because they are intrinsically motivated to uncover excellence. If experts must first choose which group to inspect, they will select groups with higher expected quality, often those with (stochastically) lower costs. Producers in these groups, in turn, are more strongly incentivized to supply quality due to the increased scrutiny: Experts' attention creates a self-fulfilling phenomenon. Moreover, when multiple experts decide which group to inspect, they naturally gravitate toward groups already receiving greater attention, as this increases the likelihood of discovering higher quality in equilibrium. As a result, endogenous attention leads to two key effects: self-fulfilling prophecies regarding group quality and expert herding in group choice. While this mechanism amplifies preexisting inequality between groups in terms of production costs, it also means that two identical populations of producers could end up in vastly different equilibria. Understanding these endogenous attention phenomena and their consequences deserves future research.

References

Ali, H. H., S. Lecocq, and M. Visser (2008). The impact of gurus: Parker grades and en primeur wine prices. *Economic Journal 118*(529), F158–F173.

Andersson, F. (2002). Pooling reputations. *International Journal of Industrial Organization 20*(5), 715–730.

Bachmann, R., G. Ehrlich, Y. Fan, D. Ruzic, and B. Leard (2023). Firms and collective reputation: a study of the volkswagen emissions scandal. *Journal of the European Economic Association 21*(2), 484–525.

Bai, J., L. Gazze, and Y. Wang (2022). Collective Reputation in Trade: Evidence from the Chinese Dairy Industry. *Review of Economics and Statistics 104*(6), 1121–1137.

Bar-Isaac, H. and S. Tadelis (2008). Seller reputation. *Foundations and Trends in Microeconomics 4 (4)*, 273–351.

Blume, L. E. (2006). The dynamics of statistical discrimination. *Economic Journal 116*(515), F480–F498.

Board, S. (2009). Monopolistic group design with peer effects. *Theoretical Economics 4*(1), 89–125.

Board, S. and M. Meyer-ter-Vehn (2013). Reputation for quality. *Econometrica 81 (6)*, 2381–2462.

Cabral, L. M. (2009). Umbrella branding with imperfect observability and moral hazard. *International Journal of Industrial Organization 27*(2), 206 – 213.

Cabral., L. M. B. (2000). Stretching firm and brand reputation. *RAND Journal of Economics 31*(4), 658–73.

Choi, J. P. (1998). Brand extension as informational leverage. *Review of Economic Studies 65*(4), 655–669.

Coate, S. and G. C. Loury (1993). Will affirmative-action policies eliminate negative stereotypes? *American Economic Review 83*(5), 1220–1240.

Drori, G. S., J. Tienari, and A. Waeraas (2015). Strategies of building and managing higher education brands. *International Studies of Management & Organization 45(2)*, 99–144.

Dubey, P. and J. Geanakoplos (2010). Grading exams: 100,99,98 ... or A,B,C? *Games and Economic Behavior 69*(1), 72–94.

Fishman, A., I. Finkelstein, A. Simhon, and N. Yacouel (2018). Collective brands. *International Journal of Industrial Organization 59*, 316–339.

Fleckinger, P. (2007). Collective reputation and market structure: Regulating the quality vs quantity trade-off. Cahier no. 2007-26, Ecole Polytechnique - CNRS.

Fleckinger, P., M. Glachant, and G. Moineville (2017). Incentives for quality in friendly and hostile informational environments. *American Economic Journal: Microeconomics 9(1)*, 242–274.

Gergaud, O., F. Livat, B. Rickard, and F. Warzynski (2017). Evaluating the net benefits of collective reputation: The case of bordeaux wine. *Food Policy 71*, 8 – 16.

Hakenes, H. and M. Peitz (2008). Umbrella branding and the provision of quality. *International Journal of Industrial Organization 26*(2), 546–556.

Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *American Economic Review 101*(6), 2590–2615.

Kim, Y.-C. and G. C. Loury (2018). Collective reputation and the dynamics of statistical discrimination. *International Economic Review 59*(1), 3–18.

Kreps, D. M. and R. Wilson (1982). Reputation and imperfect information. *Journal of Economic Theory 27*(2), 253–279.

Levin, J. (2009). The dynamics of collective reputation. *B.E. Journal of Theoretical Economics - Contributions 9*, art. 27.

Levin, J. and S. Tadelis (2005). Profit sharing and the role of professional partnerships. *Quarterly Journal of Economics 120*(1), 131–171.

MacLeod, W. B. (2007). Reputations, relationships, and contract enforcement. *Journal of Economic Literature 45(3)*, 595–628.

MacLeod, W. B. and M. Urquiola (2015). Reputation and school competition. *American Economic Review 105(11)*, 3471–3488.

Miklós-Thal, J. (2012). Linking reputations through umbrella branding. *Quantitative Marketing and Economics 10*(3), 335–374.

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

Moldovanu, B., A. Sela, and X. Shi (2007). Contests for status. *Journal of Political Economy 115*(2), 338–363.

Neeman, Z., A. Öry, and J. Yu (2019). The benefit of collective reputation. *RAND Journal of Economics 50*(4), 787–821.

Nocke, V. and R. Strausz (2023). Collective brand reputation. *Journal of Political Economy 131*(1), 1–58.

Robinson, J. (1933). *The economics of imperfect competition*. McMillan.

Saak, A. E. (2012). Collective reputation, social norms, and participation. *American Journal of Agricultural Economics 94(3)*, 763–784.

Schmalensee, R. (1981). Output and welfare implications of monopolistic third-degree price discrimination. *American Economic Review 71*(1), 242–247.

Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics 87*(3), 355–374.

Tirole, J. (1996). A theory of collective reputations (with applications to the persistence of corruption and to firm quality). *Review of Economic Studies 63*, 1–22.

Varian, H. R. (1985). Price discrimination and social welfare. *American Economic Review 75*(4), 870–875.

Winfree, J. and J. McCluskey (2005). Collective reputation and quality. *American Journal of Agricultural Economics 87*, 206–213.

Zhang, K. (2015, April). Breaking free of a stereotype: Should a domestic brand pretend to be a foreign one? *Marketing Science 34*(4), 539–554.

A Information Environments: Discussion & Extensions

The assumption of a friendly environment we have made throughout the paper deserves a substantial discussion. Additional insights, and sometimes contrasting results, appear in *neutral* or *hostile* informational environments. This section is dedicated to an extended discussion of the differences between informational environments, and presents the counterparts to our main results.

A.1 Selection Bias in Quality Revelation

In the main body of this paper the friendly informational environment comes with a level of attention specific to good quality. Our analysis has therefore relied on a particular form of selective bias: only instances of good quality are revealed. As argued already, this assumption fits well numerous examples, such as wine or education, and any environment where, more broadly, good news are more likely to surface than bad news. In other circumstances, the nature of feedback is instead hostile: only bad news tend to reach the public. Various scandals, ranging from contaminated milk to the Volkswagen emission fraud, point to a hostile environment. This is typical of scenarios with imperfect compliance. In normal circumstances, no news emerges, but crises triggered by bad news can have widespread repercussions across producers offering related products, even though their responsibility is not directly at stake. Before delving into this alternative informational environment, it is imperative to emphasize that selective news transmission is indispensable for collective reputation to impact efficiency.

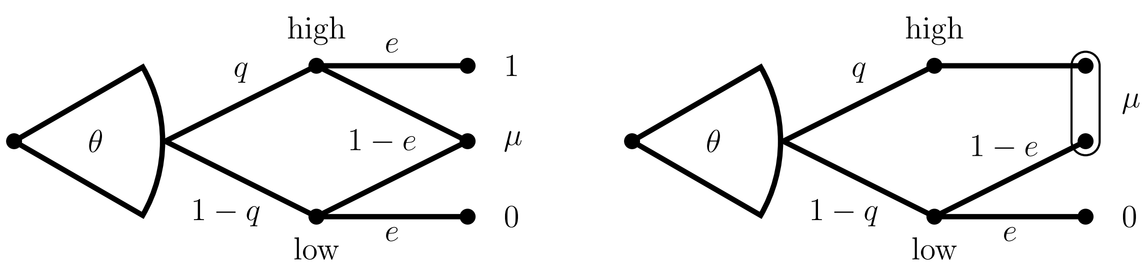

Figure 11: Neutral (left) and hostile (right) informational environment.

A neutral environment reveals quality in a non-selective way, as it features all-or-nothing monitoring. Quality is either revealed with a probability of e or remains undisclosed, as depicted on Figure 11. Then, for any μ, the incentive constraint guaranteeing that an agent with cost θ provides quality writes as follows:

$$e + (1 - e)\mu - \theta \geq e.0 + (1 - e)\mu, \tag{26}$$

which simply amounts to

$$\theta \leq e.$$

Hence, producers' incentives are independent of the collective reputation structure, including the limit case of IR. In terms of efficiency, group composition (and hence group design) is as a consequence irrelevant in a neutral informational environment.

Proposition 13 (**Irrelevance of CR in neutral informational environments**)
When the informational environment is neutral, collective reputation is irrelevant.

The proposition does not mean that a producer is indifferent to his group membership in a neutral environment: μ does enter payoffs, and everyone still seeks to belong to a group with the highest possible reputation. But in a neutral environment, reputation constitutes only a source of rents, not incentives.

A.2 Collective Reputation in a Hostile Environment

The seminal model by Tirole (1996) presents a specific bias in news selection: only bad news, such as corruption in his main example, can come to light, which makes the environment purely hostile. As previously discussed, this assumption is pivotal in generating multiple equilibria and explaining the persistence of bad equilibria in dynamic models. This insight forms the crux of Tirole's analysis, which we will not reiterate here. Instead, we aim to explore other ramifications of hostile environments for the design of collective reputation, and juxtapose them with the friendly case.

The essential difference is that incentives come from the payoff gap between no news and bad news, hence from the reputation of the group compared to a price of 0. In fact, while in a friendly environment a producer is paid *at least* his (group) reputation, here a producer is paid *at most* his (group) reputation. The corresponding incentive constraint writes as:

$$\mu - \theta \geq e.0 + (1-e)\mu, \tag{27}$$

which amounts to

$$\theta \leq e\mu.$$

The crucial implication is that reputation and incentives now go hand in hand. A higher reputation motivates the producers to not lose it by producing low quality, and conversely, a bad reputation kills incentives.

Individual Reputation

In the case of Individual Reputation, the analysis is slightly more complex than in a friendly environment, because multiple equilibria often coexist. Indeed, for any θ and any e, an equilibrium with zero effort always exists: if buyers believe that the producer chooses $q(\theta) = 0$, then $\mu(\theta) = 0$ and no incentives are created, which confirms the belief. In turn, an equilibrium with $q(\theta) = 1$ exists as soon as the corresponding belief $\mu(\theta) = 1$ indeed creates incentives, that is when $1-\theta \geq 1-e$: for all producers with costs lower than e. In addition, a mixed equilibrium exists. In such an equilibrium, a producer with type θ must be indifferent to q, given the buyers' correct belief $\mu(\theta)$. This implies $\mu^*(\theta) = \frac{\theta}{e}$, which then entails $q^*(\theta) = \frac{(1-e)\theta^*}{e(1-\theta^*)}$. Notice that in this interior equilibrium, the quality produced increases with the cost, which is clearly an undesirable feature regarding efficiency. Moreover, this equilibrium is unstable. To summarize, IR is plagued by multiple equilibria, which gives rise to an extreme form of indeterminacy: *pointwise*, the two extreme equilibria can coexist, as well then as an unstable interior equilibrium.

Collective Reputation

Turning now to the case of (Full) Collective Reputation, and following the same analysis as previously, we obtain that a perfect Bayesian equilibrium for Full Collective Reputation in a hostile environment is characterized by a threshold θ^* satisfying:

$$F(\theta^*) = \frac{(1-e)\theta^*}{e(1-\theta^*)}. \tag{28}$$

Here, too, multiple equilibria can exist. In particular, the zero-effort equilibrium (with $\theta^* = 0$) always exists. We will focus on the example of the uniform distribution in which it does not occur. In this example, Equation (28) has two solutions, but the zero-effort equilibrium is unstable: any slight perturbation triggers an upward dynamics in reputation, which in turns increases motivation of more producers, until they collectively reach the high equilibrium. Moreover, a grain of friendliness (with some positive probability of good quality revelation) or the presence of an arbitrarily small but positive mass of producers always producing high quality would destroy the zero-effort equilibrium. We can thus mostly ignore it. Hence the multiplicity issue arising with IR is resolved in this example with FCR. Clearly the resulting unique equilibrium is better than the zero-effort equilibrium in which a producer under IR can always be trapped, but it does not reach the efficiency of the best equilibrium under IR, where producers with $\theta^* < \theta \leq e$ produce additional high quality. While its relative efficiency is not guaranteed, FCR at least attenuates the indeterminacy issue.

Admission threshold

There is a simple way with CR of reaching the best outcome under IR, even sometimes uniquely: it suffices to set an admission threshold $\sigma = e$.[23] Applying the equilibrium relation to this CR structure, we obtain that there always exists an equilibrium in which producers

[23]Here it is more convenient formally to assume that the threshold is not strict, so that producers with $\theta = \sigma$ belong to group 1, but this again, is a harmless convention. Contrary to a friendly environment, a sufficiently tough standard, i.e. a low σ, prevents free-riding by the group's high-cost producers.

from group 1 all choose $q = 1$, and those in group 2 choose $q = 0$. With the uniform distribution, for instance, the admission threshold leads uniquely to the best equilibrium that can be attained under IR.[24] The rationale for the admission threshold here is however different from that in the friendly case. In a friendly environment, an admission threshold helps reducing free-riding of producers in group 2. Its purpose here is to increase incentives of producers in group 1, by letting them enjoy a motivating higher reputation. The takeaway is that a simple CR design with admission threshold can stabilize the market on the most efficient outcome attainable under IR.

Group design with transfers

Finally, group design with transfer is also markedly different from the friendly case. First of all, group stability may be harder to sustain: if a producer anticipates that he would be part of the efficient equilibrium under IR, he must then be paid at least his first-best payoff in the group. That those types below e are hard to attract in turn creates an additional issue: they are precisely those needed to motivate the others by increasing the group reputation.

When the collective organization can punish a bad product with a penalty x, it is possible to implement the first-best: the producers join the group freely, enjoy the collective reputation μ for sure when choosing high quality, and the collective organization polices them by levying the penalty if bad quality is revealed. One can easily verify that a penalty $x^* = \frac{1}{e} - 1$ does implement the first-best. However, such a solution might be problematic regarding producers' liability, since the penalty becomes arbitrarily high when information is poor.

Another solution would be to bootstrap reputation by letting the producers pay a participation fee t, which acts as a deposit recovered through a subsidy $s = t$ if no bad news occurs. Unfortunately, this scheme is plagued by exactly the same problem that it requires potentially a very high fee: implementing the first-best requires $t^* = s^* = \frac{1}{e} - 1$, and entails potentially a considerable risk as well.

[24] As discussed before, the zero-equilibrium is here (highly) unstable for the same reasons as before.

The impassable obstacle when implementing the first-best is that in a purely hostile environment, no news ever surface in equilibrium. Since all producers always receive the same price in the efficient equilibrium, based on the no news outcome, there is here no possibility to offer conditional bonuses along the equilibrium path: the market price is useless for incentives. In particular, a bonus club cannot operate. This constitute a key difference between friendly and environments for group design with transfers.

To compensate for this lack of feasible incentives along the equilibrium path, more liability is required from producers[25] to make out-of-equilibrium threats credible. In a hostile environment, reputation is both the source of revenue and incentives. Since taxing reputation reduces incentives, it cannot be used to finance bonuses unlike in friendly environment–in other words, it is never possible to pay a producer more than the group reputation. The need for upfront financing or deep pockets ex-post to create incentives is hence more pressing than in a friendly environment.

As a positive final note, it is important to bear in mind that even a tiny grain of kindness–a positive probability of identifying high quality–is sufficient to organize efficient bonus clubs or a collective retail channels. Hence all the efficiency results obtained in a friendly environment apply in general, provided the environment is not purely hostile.

B Proofs

B.1 Proof of Proposition 1

For a producer with type θ, the payoff in case of no news depends on a belief $\mu(\theta)$ specific to that type of producer. When the producer chooses a strategy $q(\theta) \in [0, 1]$, the corresponding reputation is $\mu(\theta) = \frac{(1-e)q(\theta)}{1-eq(\theta)}$ by applying Bayes' rule. For the producer to rationally choose this strategy, it must be that given $\mu(\theta)$, the first-order condition holds. Since Π is linear in q, this is akin to the indifference property in mixed strategy equilibria, i.e. indifference

[25] Recall that there exists a scheme implementing first-best with zero liability in a friendly environment.

between choosing $q = 1$ and $q = 0$, which entails $\mu(\theta) = 1 - \frac{\theta}{e}$. Combining this with the Bayesian revision yields the two equations. By the indifference property again, a producer's equilibrium payoff must equal his payoff when shirking, i.e. his reputation with probability one.

B.2 Proof of Proposition 2

Consider group i. Combining equations (12) and (13), we obtain equations (14) and (15). Note that the left-hand side of (14) increases with θ_i^*, whereas the right-hand side decreases with θ_i^*, so that there exists a unique solution. It is interior when $0 < F_i(e) < 1$, or on the boundary, either with zero quality when $F_i(e) = 0$ or full quality when $e = 1$.

B.3 Proof of Proposition 3

The first and second items come directly from inspection of (14) and (15). The third item comes from combining (13) with the second item.

The fourth item is obtained by differentiating (14) with respect to e, and rearranging to obtain: $\frac{d\theta^*}{de} = \frac{1-(1-\theta^*)F(\theta^*)}{1-eF(\theta^*)+e(1-\theta^*)f(\theta^*)} > 0$ for all $e \in (0,1)$. For point a of the fifth item, differentiating (15) with respect to e yields $\frac{d\mu^*}{de} = [-F(\theta^*) + (1-\theta^*)f(\theta^*)]\,\frac{d\theta^*}{de}$, so that the two limit cases yield $\frac{d\mu^*}{de}\Big|_{e=0} = f(0).\,\frac{d\theta^*}{de}\Big|_{e=0} \geq 0$ and $\frac{d\mu^*}{de}\Big|_{e=1} = -\,\frac{d\theta^*}{de}\Big|_{e=1} \leq 0$. For point b of the fifth item, F log-concave says f/F is decreasing. This implies that there is a threshold θ below which $[-F(\theta) + (1-\theta)f(\theta)] = \frac{F(\theta)}{(1-\theta)}\left[\frac{f(\theta)}{F(\theta)} - \frac{1}{(1-\theta)}\right]$ is positive and above which it becomes negative, because the functions f/F and $1/(1-\theta)$ are respectively decreasing and increasing, and cross in the interval $(0,1)$. Since $\frac{d\theta^*}{de} > 0$, this allows to conclude.

B.4 Proof of Proposition 4

Let us consider the sequence of collective reputation structures $(\mathcal{I}_n)_{n\in\mathbb{N}^*}$ such that for any n, the cost interval is divided in n segments of equal length. Group $i \in \{1, .., n-1\}$ consists of all producers with type $\theta \in [\frac{i-1}{n}, \frac{i}{n})$, and group n includes types $\theta \in [\frac{n-1}{n}, 1]$. For further reference, we denote by $\theta_{i,n} = \frac{i}{n}$ the group cutoffs.

Our aim is to show that the sequence of equilibria under the collective reputation structures $(\mathcal{I}_n)$ converges to the individual reputation equilibrium of Proposition 1. To be precise, we start by showing that aggregate quality converges to the same level as in the IR equilibrium, and that they do so in each of the subintervals.

Under $\mathcal{I}^n$, the cost distribution in group i is given by the truncated distribution

$$F_{i,n}(\theta) = \frac{F(\theta) - F(\theta_{i-1,n})}{F(\theta_{i,n}) - F(\theta_{i-1,n})}, \tag{29}$$

and the weight of group i is

$$\lambda_{i,n} = F(\theta_{i,n}) - F(\theta_{i-1,n}). \tag{30}$$

Let $\theta^*_{i,n}$ denote the equilibrium threshold in group i under $\mathcal{I}_n$. From Proposition 2, average quality in group i is $F_{i,n}(\theta^*_{i,n})$. Aggregating over groups, total quality is

$$Q_n = \sum_i \lambda_{i,n} F_{i,n}(\theta^*_{i,n}). \tag{31}$$

In turn, total quality in equilibrium under IR is

$$Q_{IR} = \int_0^1 q^*(\theta) f(\theta) d\theta, \tag{32}$$

which, for any n, we can also write as

$$Q_{IR} = \sum_i \int_{\theta_{i-1,n}}^{\theta_{i,n}} q^*(\theta) f(\theta) d\theta. \tag{33}$$

Since q^* is a decreasing function and $\int_{\theta_{i-1,n}}^{\theta_{i,n}} f(\theta) d\theta = \lambda_{i,n}$, it is the case that

$$\sum_i \lambda_{i,n} q^*(\theta_{i,n}) \leq Q_{IR} \leq \sum_i \lambda_{i,n} q^*(\theta_{i-1,n}). \tag{34}$$

Using now the fact that in equilibrium $F_{i,n}(\theta^*_{i,n}) = q^*(\theta^*_{i,n})$, then substracting (31),

$$\sum_i \lambda_{i,n} \left(q^*(\theta_{i,n}) - q^*(\theta^*_{i,n})\right) \leq Q_{IR} - Q_n \leq \sum_i \lambda_{i,n} \left(q^*(\theta_{i-1,n}) - q^*(\theta^*_{i,n})\right). \tag{35}$$

In addition, for all n and all i:

$$\theta_{i,n} \leq \theta^*_{i,n} \leq \theta_{i-1,n} \text{ and } |\theta_{i,n} - \theta_{i-1,n}| = \frac{1}{n}. \tag{36}$$

Since q^* is continuous over the compact interval $[0, 1]$, it is also uniformly continuous by the Heine-Cantor theorem. Hence, for any ε, there exists $N(\varepsilon)$ large enough such that for any group i in the CR structure $\mathcal{I}_n$:

$$\left|q^*(\theta_{i,n}) - q^*(\theta^*_{i,n})\right| < \varepsilon \text{ and } \left|q^*(\theta_{i-1,n}) - q^*(\theta^*_{i,n})\right| < \varepsilon \text{ for all } n > N(\varepsilon) \text{ and all } i. \tag{37}$$

In words, for a fine enough partition, quality under the sequence of CR structure converges within each group defined by the CR structure to the quality provided under IR. Since $\sum_i \lambda_{i,n} = 1$ for all n, the aggregate quality also converges as (35) holds.

Using the same method, it can also be demonstrated that the costs and welfare in the sequence of CR uniformly converge to the corresponding IR outcome. While the expression for local costs in equilibrium is slightly more intricate than the expression for aggregate quality, namely $\theta q^*(\theta)$ instead of $q^*(\theta)$, the same logical approach applies. The only added difficulty is that equilibrium costs are not monotonic, but the logic of uniform convergence is identical. The proofs are omitted, as they require only minor additional steps.

B.5 Proof of Proposition 5

For the first observation, the premise of equal quality under FCR and IR requires

$$\int_0^{\theta^*} f(\theta)d\theta = \int_0^e q^*(\theta)f(\theta)d\theta,$$

so that

$$\int_0^{\theta^*} (1 - q^*(\theta))f(\theta)d\theta = \int_{\theta^*}^{e} q^*(\theta)f(\theta)d\theta.$$

The cost difference between FCR and IR writes:

$$\begin{aligned}
C_{FCR} - C_{IR} &= \int_0^{\theta^*} \theta f(\theta)d\theta - \int_0^{e} \theta q^*(\theta)f(\theta)d\theta \\
&= \int_0^{\theta^*} \theta(1 - q^*(\theta))f(\theta)d\theta - \int_{\theta^*}^{e} \theta q^*(\theta)f(\theta)d\theta \\
&\leq \theta^* \int_0^{\theta^*} (1 - q^*(\theta))f(\theta)d\theta - \int_{\theta^*}^{e} \theta q^*(\theta)f(\theta)d\theta \\
&= \theta^* \int_{\theta^*}^{e} q^*(\theta)f(\theta)d\theta - \int_{\theta^*}^{e} \theta q^*(\theta)f(\theta)d\theta \\
&= -\int_{\theta^*}^{e} (\theta - \theta^*)q^*(\theta)f(\theta)d\theta \leq 0,
\end{aligned}$$

where the fourth line uses the premise. Moreover, since $\theta^* < e$ when $0 < e < 1$, the inequality is strict in the interior.

For the second observation, we prove the contrapositive. Assume that quality is lower under FCR than IR:

$$\int_0^{\theta^*} f(\theta)d\theta \leq \int_0^{e} q^*(\theta)f(\theta)d\theta,$$

so that

$$\int_0^{\theta^*} (1 - q^*(\theta))f(\theta)d\theta \leq \int_{\theta^*}^{e} q^*(\theta)f(\theta)d\theta.$$

Then one has

$$\begin{aligned}
\int_0^{\theta^*} \theta(1 - q^*(\theta))f(\theta)d\theta \leq \theta^* \int_0^{\theta^*} (1 - q^*(\theta))f(\theta)d\theta& \\
\leq \theta^* \int_{\theta^*}^{e} q^*(\theta)f(\theta)d\theta \leq \int_{\theta^*}^{e} \theta q^*(\theta)f(\theta)d\theta&,
\end{aligned}$$

where the middle inequality comes from the premise. Rearranging the two extreme terms yields the desired conclusion $C_{FCR} \leq C_{IR}$.

Finally, for the third observation, assume that quality is lower under FCR than IR:

$$\int_0^{\theta^*} f(\theta)d\theta \geq \int_0^{e} q^*(\theta)f(\theta)d\theta,$$

from which we obtain

$$\int_0^{\theta^*} (1-q^*(\theta))f(\theta)d\theta \geq \int_{\theta^*}^{e} q^*(\theta)f(\theta)d\theta.$$

Then,

$$\begin{aligned}\int_0^{\theta^*} (1-\theta)(1-q^*(\theta))f(\theta)d\theta &\geq (1-\theta^*)\int_0^{\theta^*} (1-q^*(\theta))f(\theta)d\theta \\ &\geq (1-\theta^*)\int_{\theta^*}^{e} q^*(\theta)f(\theta)d\theta \geq \int_{\theta^*}^{e}(1-\theta)q^*(\theta)f(\theta)d\theta,\end{aligned}$$

where the middle inequality comes from the premise. Rearranging the two extreme terms yields the desired conclusion $W_{FCR} \geq W_{IR}$.

B.6 Proof of Proposition 6

The first part is obvious: $W_{FCR}(0) = W_{IR}(0) = 0$. For the collective reputation case, using (16), direct computation yields:

$$\frac{dW_{FCR}}{de} = (1-\theta^*)f(\theta^*)\frac{d\theta^*}{de}. \tag{38}$$

Moreover, totally differentiating the equilibrium relation (14) with respect to e, one obtains:

$$\frac{d\theta^*}{de} = \frac{1-(1-\theta^*)F(\theta^*)}{1+e\left((1-\theta^*)f(\theta^*)-F(\theta^*)\right)}. \tag{39}$$

Since $\lim_{e\to 0}\theta^* = 0$ and f is bounded because F is regular, $\lim_{e\to 0}\frac{d\theta^*}{de} = 1$, so that:

$$\lim_{e\to 0}\frac{dW_{FCR}(e)}{de} = f(0). \tag{40}$$

For the Individual Reputation case, we obtain using (11):

$$\frac{dW_{IR}}{de} = \frac{1}{e^2}\int_0^e \theta f(\theta)d\theta. \tag{41}$$

Now, let

$$m(e) \equiv \inf\{\operatorname*{argmax}_{\theta\in[0,e]} f(\theta)\}. \tag{42}$$

Note that $m(e)$ is well defined since f is bounded. It holds that:

$$\frac{dW_{IR}}{de} \leq \frac{1}{e^2}\int_0^e \theta f(m(e))d\theta = \frac{1}{2}f(m(e)). \tag{43}$$

Given the definition of m, $m(0) = 0$, so that:

$$\lim_{e\to 0}\frac{dW_{IR}}{de} \leq \frac{f(0)}{2}, \tag{44}$$

which allows to conclude by comparing the two welfare limits. In addition $f(0) > 0$ when F is regular, so that FCR then strictly dominates when e is low enough.

B.7 Proof of Proposition 7

The first part is obvious: $W_{FCR}(1) = W_{IR}(1) = W_{FB}$. For the collective reputation case, we had obtained in the proof of Proposition 6:

$$\frac{dW_{FCR}}{de} = (1-\theta^*)f(\theta^*)\frac{d\theta^*}{de},$$

and

$$\frac{d\theta^*}{de} = \frac{1-(1-\theta^*)F(\theta^*)}{1+e\left((1-\theta^*)f(\theta^*)-F(\theta^*)\right)}.$$

By using the variable change $h(e) = 1-\theta^*$, we can rewrite the derivative of the welfare as follows:

$$\frac{dW_{FCR}}{de} = \frac{h(e)f(1-h(e))(1-h(e)F(1-h(e))}{1+e(h(e)f(1-h(e))-F(1-h(e))} \equiv \frac{N(e)}{D(e)}.$$

Since $\lim_{e \to 1} \frac{d\theta^*}{de} = +\infty$, h does not admit a Taylor expansion near $e = 1$, hence we cannot use a compound Taylor expansion to obtain the limit. However, note that $h(1) = 0$, so that we can write the Taylor expansion of the numerator when $e \to 1$ as:

$$N(e) \underset{e \to 1}{\sim} h(e)f(1 - h(e)) + O(h(e)^2).$$

For the denominator, we use the first-order Taylor expansion of $F(1 - h(e))$, which is well defined for a regular F, and obtain:

$$\begin{aligned} D(e) &\underset{e \to 1}{\sim} 1 + h(e)f(1 - h(e)) - (F(1) - h(e)f(1 - h(e)) + O(h(e)^2) \\ &\underset{e \to 1}{\sim} 2h(e)f(1 - h(e)) + O(h(e)^2). \end{aligned}$$

We can therefore conclude that:

$$\lim_{e \to 1} \frac{dW_{FCR}}{de} = \frac{1}{2}.$$

Consider now the derivative of the welfare under individual reputation:

$$\frac{dW_{IR}}{de} = \frac{1}{e^2} \int_0^e \theta f(\theta) d\theta,$$

so that:

$$\lim_{e \to 1} \frac{dW_{IR}}{de} = \mathbb{E}[\theta].$$

Since the welfare when $e = 1$ is the same under IR and FCR, IR dominates FCR in a neighborhood of $e = 1$ if and only if the derivative of the welfare under IR is lower. This yields the second equation of the Proposition.

B.8 Proof of Proposition 8

he welfare difference is given by:

$$
\begin{aligned}
W_{FCR} - W_{IR} &= (1-\theta^*)F(\theta^*) + \int_0^{\theta^*} F(\theta)d\theta - \frac{1}{e}\int_0^{e} F(\theta)d\theta \\
&= (1-\theta^*)\theta^* + \int_0^{\theta^*} \theta d\theta - \frac{1}{e}\int_0^{e} \theta d\theta \\
&= (1-\theta^*)\theta^* + \frac{(\theta^*)^2}{2} - \frac{e}{2} \\
&= \frac{1}{2}(1-e)\theta^*(1-\theta^*),
\end{aligned}
$$

where we use (14) to obtain that $e - \theta^* = e\theta^*(1-\theta^*)$m which proves the proposition.

B.9 Proof of Proposition 9

For any collective reputation structure, consider the lowest cost type $\underline{\theta}_i$ in some group i that produces quality. The payoff of this type is $U_{IR}(\underline{\theta}_i) = \mu(\underline{\theta}_i) = 1 - \frac{\underline{\theta}_i}{e}$ under individual reputation and $U_{CR}(\underline{\theta}_i) = \mu_i^* + \theta_i^* - \underline{\theta}_i = 1 - \frac{\theta_i^*}{e} + \theta_i^* - \underline{\theta}_i$ under collective reputation. Therefore $U_{IR}(\underline{\theta}_i) - U_{CR}(\underline{\theta}_i) = (\theta_i^* - \underline{\theta}_i)(\frac{1}{e} - 1) > 0$.

B.10 Proof of Proposition 10

Following the proof of Proposition 2, equation (18) has a unique solution θ_b^*. It is immediate to check that $\theta_b^* = 1$ satisfies (18) when $b = \frac{1}{e}$. As a consequence, $\mu_b^* = 1$. A producer's profit is then $1 - \theta$. The participation constraint (19) can now be written as $\mu^*(\theta) + \theta \leq 1$, or by substituting the individual reputation obtained in Proposition 1: $1 - \theta\left(\frac{1}{e} - 1\right) \leq 1$, which always holds. Note finally that this scheme balances the budget exactly. Indeed, a fee of 1 is levied on all producers, generating a budget of 1, while the bonus $\frac{1}{e}$ is paid with probability e, for an aggregate cost of 1.

B.11 Proof of Proposition 11

Participating in the group amounts here to obtaining a lottery ticket that is worthless without effort, but worth 0 with probability $(1-e)$ and $1/e$ with probability e when providing quality. In total, participation with effort hence yields the first-best payoff $1-\theta$, which beats the alternative of not participating, and efficiency is maximized. The scheme is clearly budget-balanced since all prices on the market are 1 in equilibrium, which covers the total amount redistributed.

B.12 Proof of Proposition 12

The welfare with cutoff σ can be written as:

$$W(\sigma) = \int_0^{\theta_1^*(\sigma)} (1-\theta)f(\theta)d\theta + \int_\sigma^{\theta_2^*(\sigma)} (1-\theta)f(\theta)d\theta.$$

Note that $W(0) = W(1) = W_{FCR}$. The derivative of welfare with respect to σ is

$$\frac{dW}{d\sigma} = (1-\theta_1^*(\sigma))f(\theta_1^*(\sigma))\frac{d\theta_1^*}{d\sigma} + (1-\theta_2^*(\sigma))f(\theta_2^*(\sigma))\frac{d\theta_2^*}{d\sigma} - (1-\sigma)f(\sigma).$$

We hence have

$$\left.\frac{dW}{d\sigma}\right|_{\sigma=0} = f(0)(\left.\frac{d\theta_1^*}{d\sigma}\right|_{\sigma=0} - 1) + (1-\theta_2^*(0))f(\theta_2^*(0))\left.\frac{d\theta_2^*}{d\sigma}\right|_{\sigma=0}.$$

Applying Proposition 2 to the two groups, and differentiating (14) with respect to σ in both groups, we obtain:

$$\frac{d\theta_1^*}{d\sigma} = \frac{(e-\theta_1^*)f(\sigma)}{F(\sigma) - eF(\theta_1^*) + e(1-\theta_1^*)f(\theta_1^*)}$$

and

$$\frac{d\theta_2^*}{d\sigma} = \frac{(1-e)\theta_2^* f(\sigma)}{1-(1-e)F(\sigma) + e(1-\theta_2^*)f(\theta_2^*) - eF(\theta_2^*)}.$$

Hence, since $f > 0$ for a regular F, we have

$$\left.\frac{d\theta_2^*}{d\sigma}\right|_{\sigma=0} = \frac{(1-e)\theta_2^* f(0)}{1 + e(1-\theta_2^*)f(\theta_2^*) - eF(\theta_2^*)} > 0,$$

so that the second term in the welfare derivative is positive. Moreover, when $\sigma = 0$, $\theta_1^* = 0$, hence $\frac{d\theta_1^*}{d\sigma}|_{\sigma=0} = \frac{ef(0)}{ef(0)} = 1$, so that the first term in the welfare is zero. As a result, $\frac{dW}{d\sigma}|_{\sigma=0} > 0$, and there exists an interior cutoff σ strictly improving on the single group.

C Proofs for the Uniform Distribution

In this appendix we obtain and gather all explicit expressions pertaining to the case of the uniform distribution.

Individual Reputation With individual reputation, the quality produced by type θ is $q^*(\theta) = \frac{e-\theta}{1-\theta}$ irrespective of the distribution. From Equation (11), the welfare under individual reputation is:

$$\begin{aligned} W_{IR}(e) &= \frac{1}{e}\int_0^e d\theta \\ &= \frac{1}{2}e. \end{aligned}$$

The total quality produced is in turn:

$$\begin{aligned}
Q_{IR}(e) &= \int_0^e \frac{e-\theta}{e(1-\theta)}\theta d\theta \\
&= \frac{1}{e}\left([-(e-\theta)\ln(1-\theta)]_0^e - \int_0^e \ln(1-\theta)d\theta\right) \\
&= 0 - \frac{1}{e}\int_0^e \ln(1-\theta)d\theta \\
&= \frac{1}{e}\left[(1-\theta)(\ln(1-\theta)-1)\right]_0^e \\
&= 1 + \frac{1-e}{e}\ln(1-e).
\end{aligned}$$

Collective Reputation With collective reputation, we must first determine the equilibrium threshold θ^*. Applying Equation (14), the threshold θ^* is the solution to:

$$e\theta^*(1-\theta^*) = e - \theta^*,$$

which has a unique root in the unit interval:

$$\theta^* = \frac{1}{2} + \frac{1-\sqrt{(1-e)(1+3e)}}{2e}.$$

The welfare is then obtained from Equation (16):

$$\begin{aligned}
W_{CR}(e) &= \int_0^{\theta^*}(1-\theta)d\theta \\
&= \theta^*(1-\frac{\theta^*}{2}) \\
&= \frac{(1+e-\sqrt{(1-e)(1+3e)})(1-3e-\sqrt{(1-e)(1+3e)})}{8e^2}.
\end{aligned}$$

The quality produced is $Q_{CR}(e) = F(\theta^*)$, hence:

$$Q_{CR}(e) = \frac{1}{2} + \frac{1-\sqrt{(1-e)(1+3e)}}{2e}.$$

These expressions allow to plot the figures presented in the main body of the paper and are used in the proofs of the propositions pertaining to the uniform distribution.

Admission Threshold We gather here the explicit expressions for the analysis of admission thresholds. Applying the equilibrium characterization of Proposition 2, the equilibrium thresholds in groups 1 and 2 are obtained by solving the two corresponding second degree polynomials, each having a unique root in the relevant intervals. They are given by

$$\theta_1^* = \frac{1}{2} + \frac{\sigma - \sqrt{(\sigma+e)^2 - 4e^2\sigma}}{2e}, \tag{45}$$

and

$$\theta_2^* = \max\left\{\frac{1}{2} + \frac{1-(1-e)\sigma - \sqrt{(1-e)(1-\sigma)(1-\sigma+e(3+\sigma))}}{2e}, \sigma\right\}. \tag{46}$$

Using the distributions in Equation (21) and plugging in the expression of welfare (23), we get

$$W(\sigma) = \theta_1^*(1-\frac{\theta_1^*}{2}) + (\theta_2^* - \sigma)(1 - \frac{\sigma+\theta_2^*}{2}). \tag{47}$$

Finally, the aggregate tuition for a private school with the uniform distribution is obtained from Equation (25), as

$$T(\sigma) = \begin{cases} \frac{1-e}{e}(\theta_2^* - \theta_1^*)\sigma, & \text{if } \sigma < e \\ (1-e)(1-\frac{\theta_1^*}{e})\sigma, & \text{if } \sigma \geq e. \end{cases} \tag{48}$$

These expressions allow to obtain numerically the efficient and tuition maximizing thresholds, and to plot the figures in Section 5.