The Incentive Properties of Collective Reputation

Pierre Fleckinger* Wanda Mimra† Angelo Zago‡

November 2018

PRELIMINARY DRAFT–PLEASE DO NOT CIRCULATE

Abstract

Collective Reputation is often viewed as creating free-riding and impeding quality provision. It is however a widespread and often deliberate choice of producers. We provide new explanations for this based on group incentives. Heterogeneous producers whose costs are imperfectly known need to provide effort to produce high quality. The demand side a priori does not observe the true quality, but high quality can be detected with some probability, reflecting e.g. expert inspections, awards, labeling and the regulatory framework. Unidentified products are otherwise pooled together according to the collective reputation structure, i.e. grouping of producers. In the unique equilibrium, each group is subject to internal free-riding by their higher-cost members. We find however that grouping producers can also increase incentives and yield higher quality and welfare than individual reputation, because free-riding under collective reputation might be less severe than own-reputation milking under individual reputations. We also show that admission thresholds with a small elite always improves upon full collective reputation. Despite potentially higher producers' surplus, any group with collective reputation however unravels in absence of transfers. Nevertheless, we exhibit simple type-independent and budget-balanced contracts under collective reputation that implement the first best.

JEL classification: D82; D47; D71; L15.

Keywords: Collective Reputation, Quality, Incentives.

*MINES ParisTech and Paris School of Economics. Email: pierre.fleckinger@mines-paristech.fr.
†ETH Zurich. Email: wmimra@ethz.ch.
‡University of Verona. Email: angelo.zago@univr.it.

Contents

List of Figures

1 Introduction

Collective reputation is an important institution for the functioning of many markets otherwise plagued by asymmetric information about product quality. It is quite common in manufacturing industries–think about Swiss watches, Italian clothing, German cars–where country of origin may be evocative of the average quality of products. It may be an important institution for service sectors as well, for instance in education, where school reputation is the source of perhaps the most important peer-effect. School reputation affects the choices of students, families, and school administrators[1] and it has been long recognized that graduates use their school of origin to signal ability, endogenously leading to stratification[2]. In agricultural markets collective reputations have been intensely supported and regulated as well, like for instance for many european cheeses (e.g., Camembert, Comté) or wines (e.g., Champagne, Bordeaux, Chianti), just to name a few examples.[3]

Information on those markets is provided by experts. On the wine market, experts are rather of the friendly kind, that tend to report only news that do not hurt. A selectively positive bias is indeed documented, whereby more good news surface. A typical illustration is that a wine trade at a higher price when it has received a Parker grade (Ali et al., 2008). This selective bias is easily explained by the presence of incumbent experts protecting their market influence, or experts being paid by producers, or by selective samples chosen by the producer for evaluation.

In education, friendliness of feedback seems to be a general rule. Prominently, schools delivers degrees, thereby providing students with a (positive) certification. Professors are in a vast majority rather well intentioned towards their students. In private education systems, students are clients who prefer not to receive a bad news ex-post; this fact may provide an economic explanation for a friendly bias. If teaching styles might differ at private and free public universities, a good approximation is still that they both constitute friendly environments. Of course, in other cases, for instance related to corruption of fraud, Collective

[1]See, e.g., Drori et al. (2015)).
[2]See MacLeod and Urquiola (2015).
[3]See Zago (2015)... broad overview.

Reputation is sustained through other kinds of feedbacks. While news from Professors is usually good, news from Police is usually bad. This latter situation corresponds to a hostile informational environment. Tirole (1996) classic analysis of persisting corrupt reputation relies on a hostile environment. In contrast, we emphasize here friendly environments and borrow the information structure from Fleckinger et al. (2017). Whether a particular environment prevails depends on the application. Our analysis focuses on the friendly case, but we also provide a complete discussion of the alternative environments.

Our model of collective reputation features moral hazard and adverse selection. This allows to investigate the incentives faced by firms in providing quality under alternative collective reputation structures. A continuum of heterogeneous producers need to exert effort to produce quality and differ in their cost. The demand side a priori does not observe the true quality; however, high quality is detected as such by the market with some probability through a given technology, reflecting for instance the performing of inspections, the provision of experts' recommendation, the winning of awards, and the like. Products (or services) not detected as of high quality are pooled together according to the collective reputation structure, i.e., grouping or partitioning of producers.

An important reference case corresponds to individually identified producers. Individual reputation is then tailored to the characteristic of each given producer. In that case, only the moral hazard problem still has a bite. 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 actually provide quality to get a premium over the base reputation are weaker than if this reputation is low. As a result, in the unique equilibrium, a producer with a positive cost will never produce enough quality.

When reputation is collective, heterogeneous producers are grouped together, and the market does not identify a given producer individually. Similarly to the individual reputation case, a high reputation hinders incentives. However, a producer alone does not influence the reputation of the group he belongs to, and incentives hence differ.[4] Any collective reputation

[4]By assuming a continuum of producers, our model allows to abstract from size effect. This also implies

structures yield a unique equilibrium, which is characterized by a cost threshold in each group, such that only members with costs below the threshold provide full quality, while the others fully free-ride. But while in the individual case reputation milking, i.e. free-riding on one's own reputation, is a pure efficiency loss, free-riding in the group decreases the reputation and actually creates in turn incentives for the other group members. As a result, we find that collective reputation can yield higher quality than individual reputation. Moreover, the form of free-riding under Collective Reputation is associated with an efficient cost allocation for the average quality produced. These effects in combination explain that Collective Reputation attenuates reputation-milking and can even yield higher welfare than Individual Reputation.

We study different collective reputation structures from an efficiency point of view, then investigate whether it is possible to design a stable collective reputation structure. We first study admission standards and show that full collective reputation can always be improved upon by an admission threshold based on costs. The role of an admission threshold is to insulate less efficient producers from more efficient ones in order to increase the incentives of the former without reducing too much those of the latter. The logic can be iterated at the bottom of the cost distribution, showing that a form of elitism efficiently increases the incentives of producers outside of the top group and welfare.

Unfortunately, despite aggregate producers' profit (and welfare) being potentially higher under collective reputation, this is not the case at the individual level: lowest cost producers in a group would always prefer to secede and adopt individual reputations, so that any group with collective reputation without transfers or mandatory participation will tend to unravel. Unlike the case of pure adverse selection, such unraveling is hence often inefficient and must be prevented.

Designing a stable, budget-balance Collective Reputation organization with transfers proves to be a fruitful exercise. We propose two simple type-independent and budget-

that no individual producer perceives an influence on the reputation of the group. The important feature is that reputation is less sensitive to a given producer's choice of quality under Collective Reputation than under Individual Reputation. Collective Reputation creates free-riding, but it is in part compensated because this means a reputation loss that in turns creates incentives for the group. This fundamental mechanism does not depend on the assumption of a continuum instead of an integer number of agents.

balanced transfer schemes based solely market price that implement the first-best. In the first scheme, that we label "bonus club", producers pay ex-ante a membership fee, and the funds are then use to provide bonuses on top of market prices. With a well-calibrated transfer scheme, this allows to increase incentives to the efficient level. An alternative is a "common retailer" organization, in which producers sell through a collective organization that allows to distort prices paid back to the producer to generate efficient incentives. The advantage is that producers do not need initial capital to implement this second solution. This shows how a collective organization can leverage the value created by organizing Collective Reputation to restore first-best incentives.

Finally, given the importance of the information structure in collective reputation phenomena, we extensively discuss the alternative informational feedbacks under which producers could operate. In a hostile environment, we identify a variety of other benefits of Collective Reputation compared to Individual Reputation, most notably equilibrium stabilization.

Applications

While our model has broad applicability, we illustrate it extensively in the context of the wine sector, an unsurpassed example for Collective Reputation issues discussed both at the industry and policy-making levels. We show for instance how collective reputation may increase the overall incentives for quality investments, thus explaining why some policy-makers are giving them such an important role. In the EU, for instance, Geographical Indications (GIs) are becoming so important that at some point they constituted a major obstacle in the TTIP trade agreements with the US.[5] Given that international trade is

[5]See e.g. Breteau and Audureau (2016), Tafta: pourquoi les Etats-Unis peuvent produire mozzarella, chablis ou champagne, Le Monde, February 19th. Collective reputations were probably first officially recognized and regulated in Europe, but have been (and still are) spreading all over. In the USA, American Viticultural Areas regulation has been in use since 1978 (Napa Valley was recognized in 1981). New Zealand has decided to bring into force the so called "GI Act" with a registration regulation, to establish a sytem for ackowledging and registering Geographical Indications (GIs) from abroad and from NZ. More recently, China's emerging Ningxia wine region has announced plans for its first winery classification based on a Bordeaux-like system, see Wu (2016), Ningxia announces wine classification system, Decanter China, February 16th.

increasingly based on quality competition (see, e.g., Baldwin and Harrigan, 2011; Crozet et al., 2012), this seems an understandable concern.

The assumption of a friendly environment that we use as the lead case is an accurate one for the wine industry. Experts inspect a fraction of the wines, and reports in their guides those that they deem of good quality, but they do not publicize bad quality. Similarly, expert buyers such as upscale hotels and restaurants buy themselves good wines on a primary market, and the remaining wines are sold on the general market for non-experts, who however benefit from the overall quality increase induced by expert buyers scrutiny.

Moreover, our model provides insights on practical design of optimal collective reputation structure. If a group splits, incentives for low-cost producers may decrease, while those for high-cost producers increase to the extent that free-riding is reduced. We show that there exists an optimal admission threshold to the first group such that welfare is greater with two subgroups than with one single group. This implies that group information often should be finer at the top, i.e. some degree of collective reputations and elitism is optimal. This is perfectly illustrated by the system in Bourgogne, which has a hierarchical classification where all 'climats' are identified and those at the top more finely defined.[6] However, while splitting into finer partitions may be sometimes good for quality incentives and economic welfare, not all splits are welfare-improving: a badly chosen admission standard can even destroy value. When is the case, a merger of reputations will then increase welfare.

While we argue that in some circumstances collective reputation may be welfare increasing, the lowest cost producers in any group producing quality would always prefer to secede. This implies that if participation is voluntary, we might expect unraveling from the top. The attempts of top producers to leave collective reputations have been duly reported, for instance, for Rijoa wines in Spain, Amarone wines in Italy,[7] Loire wines in France.[8,9]

[6] An interesting case matching our predictions is Montagny, a village in Bourgogne that has recently opted for a more 'elitist' collective structure, see Anson (2016), Tasting the Burgundy climats, Decanter, March 3rd.

[7] See e.g. Guerrini (2013), Amarone, sui vigneti in pianura guerra aperta, Corriere di Verona, May 7th.

[8] See e.g. Anson (2013), Montlouis-sur-Loire cuts itself loose from InterLoire, Decanter, January 9th.

[9] Stability is a recurring issue in other agricultural markets as well, as the case ot the US Milk Marketing Order illustrates (Crespi and Marette, 2002).

When investigating how to sustain collective reputations, we suggest simple schemes that allow to redistribute the baseline reputation to increase the payoff for identified quality. This is much in line with the way producers' organization in the European wine industry operate. These collective organizations are maintaining and taxing the group reputation, and target their promotion efforts to the best wines. While seemingly anti-redistributive, such a policy does create appropriate incentives, and increases in turn the reputation of the group from which all producers benefit through higher market prices. Having a high reputation can therefore be efficiently leveraged under collective reputation, while it constitute on the contrary a limit under individual reputation.

Related literature

While the contributions on individual reputation are very numerous (for a survey see, e.g., Bar-Isaac and Tadelis, 2008), those on collective reputation are still few, even though their number is growing. Collective reputation is associated with various ideas and models in the literature.

The seminal paper on collective reputation by Tirole (1996) builds on the classic game theoretic approach to reputation (*à la* Kreps and Wilson, 1982), where agents are of different 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 where high quality equilibria can be sustained (or not). One of the main results is the multiplicity of equilibria that potentially leads to long-lasting low reputation equilibria.[10] As discussed in Fleckinger et al. (2017), multiplicity is a consequence of the particular informational environment considered in Tirole (1996). It is 'hostile' in the sense that news consists of signals on low quality.[11] Multiplicity of equilibria in this informational

[10]Blume (2006) and Levin (2009) consider stochastic versions of Tirole (1996)'s collective reputation model. In such a framework, Levin (2009) shows why moving from one steady-state to another may be gradual, and why small policy changes may not favorably shift behavior following a history of poor outcomes. Kim and Loury (2018) pursue this line of inquiry by extending the original static setting of Coate and Loury (1993) to a dynamic version. At any rate, our main line of inquiry around group design differ substantially from the issues analyzed in the literature on Statistical Discrimination, where groups are given.

[11]It is a 'bad news case', in the wording of Board and Meyer-ter-Vehn (2013).

environment stems from quality incentives and beliefs going hand-in-hand: If the market believes quality is high, quality needs to be kept high in order not to send a low quality signal. If the market believes quality is low, then incentives are low, since the signal does not contradict the market belief. In this paper, we dedicate particular attention to the diametrically opposed informational environment: A friendly environment in which only high quality can be revealed. This informational environment naturally fits many applications, such as our leading example of the wine market, and more generally the examples given above, and enables unambiguous comparison between different collective reputation structure due to uniqueness of equilibrium–a distinctive feature of our analysis.

Modeling Collective Reputation has taken different forms in the literature. While Tirole (1996), Coate and Loury (1993) and Fleckinger et al. (2017) feature a continuum of agents, other papers consider atomic agents. 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 role; their focus is rather on the free-riding cost of Collective Reputation. Fishman et al. (2018) and Neeman et al. (2016) highlight on the contrary some benefits of Collective Reputation in games with granular individual influence. In particular, rather than generally characterize equilibrium, they find conditions under which Collective Reputation can implement the efficient equilibrium while Individual Reputation cannot. The advantage of our model is that we can conveniently study all equilibria, not only the ones with efficient investment. It also enables to discuss Collective Reputation design in more dimensions, in particular on the very nature of producers' grouping and collective organization that can increase incentives.

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)'s contribution lies in understanding the profit-maximizing price strategy in the design of groups. The optimal way of dividing players into categories, ex-ante or ex-post, in order to provide incentives is also adressed 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 the 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.

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). The question of comparing the incentive effects of different collective reputation structures has so far not been studied to our knowledge.

The remainder of the paper is organized as follows. In Section 2, we present our model of collective reputation, present the cases of Individual Reputation and study in detail equilibrium under Collectice Reputation. In Section 3, we compare the two arrangements, and study admission threshold. In Section 4, we highlight the scope for sustaining collective reputation and increasing incentives by internal group schemes. In Section 5, we discuss extensively our central assumption of a friendly environment and provide an analysis of the model under alternative assumptions. The last Section concludes.

2 A Model of Collective Reputation

2.1 Fundamentals

The canonical model features a unit mass of heterogeneous producers, characterized by a cost parameter $\theta \in [0, 1]$ of providing quality. This cost pertains to production technology, individual skill 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 the market knows about the characteristic θ of a given producer. In the baseline case, that we refer to as Individual Reputation, the market identifies each producer individually, so that the cost is public information. We now introduce our formal notion of Collective Reputation.

Collective Reputation Structures.

Under Collective Reputation (CR), producers are organized into groups such that the market only knows which group a producer belongs to, as well as the distribution of cost types in that group. A collective reputation structure allocates producers into a collection of disjoint groups indexed by $i \in \mathcal{G}$. We restrict attention to a countable number of groups, and for convenience, work directly with the (continuous) distributions over θ within groups. Thus, any countable collection of continuous distributions and associated weights $\{\lambda_i, F_i\}_{i \in \mathcal{G}}$ defines an admissible reputation structure when it satisfies:

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

This definition is in a sense inspired by an information design approach, since this identity is nothing else than Bayes' rule.[12] In economic terms, λ_i represents the relative weight of group i in the population of producers, and F_i is the distribution of types conditional on the producer belonging to group i. Importantly, the collective reputation structure is common knowledge but consumers only observe group membership, not the type of a producer.

Some polar collective reputation structures are worth emphasizing. We say that there is Full Collective Reputation (FCR) if there is a single group. We will later also consider the specific case of an admission threshold where producers are split in two groups according to a threshold type σ. Furthermore, we say that the reputation structure is horizontal when producers are not grouped according to θ. Note that given the assumption of a continuum of producers, this is then equivalent to replicating FCR, since the distribution is identically equal to F in all groups.

Finally, note that the Individual Reputation (IR) structure where θ is observed by consumers does not enter our definition because of an uncountable number of groups. However it may be naturally thought of as a limit where each producer is alone in his group. This will

[12] It is hence quite general, despite the countability assumption, which we essentially make for technical ease. It includes for instance Borelian structures *à la* Board (2009) in which all producers of a given type must be in the same group.

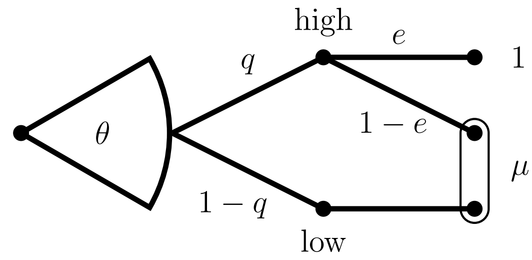

Figure 1: Information structure.

constitutes our reference case.

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 notation 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 $c(q, \theta) = \theta q$.

On the demand side, the willingness to pay is normalized to 1 for high quality and 0 for low quality.[13] Hence in general the willingness to pay is the expectation of q conditional on the available information. The demand side of the market a priori does not observe the true quality of a good from a given producer. However, in the market there are experts who provide (some) information on quality through ratings, recommendations, awards etc. We assume that high quality is publicly recognized as such with probability e, which captures the intensity of attention of experts. If quality is low, no public information is provided–the good can neither win an award or be distinguished, nor be singled out as low quality. Hence there are two types of goods available on the market: the ones which have been identified as high quality, and the unidentified ones. We assume that a product is traded at its expected quality, i.e. producers are able to charge fully the willingness to pay of the consumers, which assumes away competition. An identified high quality good is thus traded at a high price of $p = 1$. An unidentified good is traded on the basis of the *reputation* μ of the producer (or of the group he belongs to), which is the belief that the good is of high quality when

[13]It is relatively easy to extend the analysis to negative valuations for low quality, see the Appendix in Fleckinger et al. (2017).

it has not received a positive review from experts.[14] The corresponding price is therefore $p(\mu) = \mu.1 + (1-\mu).0 = \mu$, and all products are traded. For a given reputation μ, the expected payoff of a producer is then given by:

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

The actual quality and reputation that emerge will result from the architecture of collective reputation, the incentives of producers and the information structure on quality. We adopt the notion of perfect Bayesian equilibrium where each (type of) producer chooses his best reply to the market's belief, and the non-informed buyers' beliefs are consistent with the distribution of quality offered. Observe that, 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 market's belief for this group, and therefore producer incentives, are independent of the collective reputation structure for all producers not belonging to the group. Since we can analyze groups separately, we will sometimes slightly abuse wording for simplicity and speak of the *group equilibrium* when characterizing equilibrium behavior in a given group in the perfect Bayesian equilibrium of the market. Welfare in the market is given by quality minus costs. Hence, for a given collective reputation structure, we have:

$$W_{CR} = \sum_{i\in\mathcal{G}} \lambda_i \int_0^1 (1-\theta) q_i(\theta) f_i(\theta) d\theta, \tag{2}$$

where $q_i(\theta)$ is the quality chosen by a producer of quality θ in group i. The implicit assumption that all producers with the same θ in a group choose the same quality is without loss of generality as we shall see in the equilibrium analysis. Note that the welfare for the case of IR is defined analogously by integration over the individual strategy profile.

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

[14] Hence the term *reputation*, slightly abusively, refers only to unidentified product. Identified high quality alway have of course the best reputation, equal to 1. Note also that *reputation* refers here to expected quality, as in Board and Meyer-ter-Vehn (2013), and not to the type of buyer, as in more traditional reputation models.

lower than the social value of quality, yielding welfare level:

$$W_{FB} = \int_0^1 (1-\theta) f(\theta) d\theta. \\ = \int_0^1 F(\theta) d\theta. \tag{3}$$

As a note, one can easily see that given imperfect information, incentives are gonna be limited, and the First-Best level is not attainable

Incentives for Quality.

We are now in position to study equilibrium under various reputation arrangements.

2.2 Equilibrium under Individual Reputation

Individual reputation, where each producer is alone in his group, implies that each producer's θ is known by the market. 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 $\mu(\theta)$, a producer chooses $q(\theta)$ to maximize his payoff given by (1), so that:

$$q(\theta) = \begin{cases} 1 & \text{if } \theta < e(1-\mu(\theta)) \\ [0,1] & \text{if } \theta = e(1-\mu(\theta)) \\ 0 & \text{if } \theta > e(1-\mu(\theta)) \end{cases} \tag{4}$$

Moreover, the reputation $\mu(\theta)$ is consistent when:

$$\mu(\theta) = \frac{(1-e)q(\theta)}{(1-e)q(\theta) + (1-q(\theta))} \tag{5}$$

These two conditions together define a perfect Bayesian equilibrium. As seen from (4), incentives are decreasing with reputation, in particular they are zero when $\mu(\theta) = 1$. It is easy to see that mixing must be involved: producing quality with probability 1 cannot happen if the cost is positive. Indeed, if it were the case, then in equilibrium the market would attribute the best reputation to such a producer, but that would in turn kill incentives. Moreover, if the cost is too high, the producer will never produce 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

$$q^*(\theta) = Max\{\frac{e-\theta}{e(1-\theta)}, 0\}, \tag{6}$$

and the reputation of a producer of type θ is

$$\mu^*(\theta) = (1-\theta)q^*(\theta). \tag{7}$$

Proof. 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 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. ■

This equilibrium has the expected properties: More information increases incentives, higher cost producers chooses lower quality and have lower reputation. 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:

$$W_{IR}(e) = \frac{1}{e}\int_0^e F(\theta)d\theta \tag{8}$$

The Welfare under IR is strictly below the First-Best when information is imperfect. Producers are trapped in a situation where their own reputation is self-defeating: the reputation

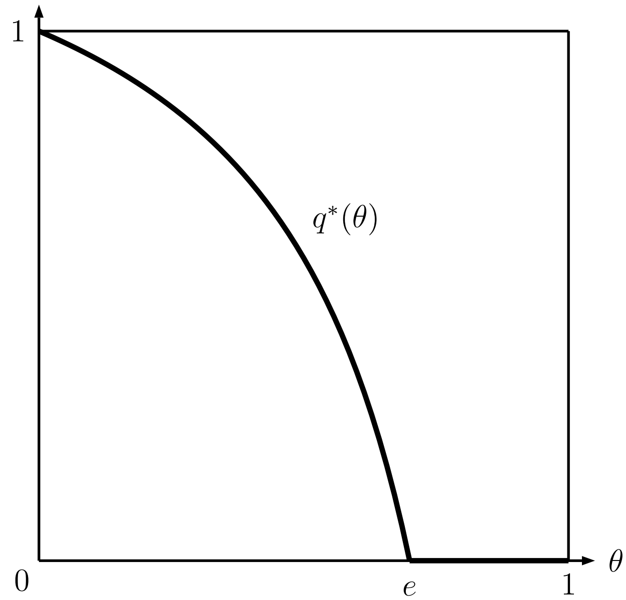

Figure 2: Equilibrium quality under Individual Reputation.

milking temptation is too high to create powerful incentive.

2.3 Equilibrium under Collective Reputation

We study now equilibrium under the Collective Reputation structures defined before. 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. In terms of incentives, a producer chooses $q = 1$ whenever $\Pi(1, \theta) \geq \Pi(0, \theta)$, which translates into $\theta \leq e(1 - \mu_i)$. An almost immediate consequence of these incentives considerations is that in all groups, equilibria have a cutoff structure: they will be characterized by a cost threshold θ_i^* below which producers of group i choose high quality and above which producers of group i choose low quality, that is:

$$\theta_i^* = e(1 - \mu_i^*). \tag{9}$$

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

$$\mu_i^* = \frac{(1-e)F_i(\theta_i^*)}{(1-e)F_i(\theta_i^*) + (1 - F_i(\theta_i^*))}. \tag{10}$$

Equations (9) and (10) characterize the equilibrium, as described in the next proposition.

Proposition 2 (**Equilibrium under Collective Reputation**)
For any admissible collective reputation structure, there is a unique equilibrium, such that for all i, there exists a threshold θ_i^* such that in group i, a producer of type $\theta \leq \theta_i^*$ chooses $q = 1$, and $q = 0$ otherwise. Average quality in group i is

$$F_i(\theta_i^*) = \frac{e - \theta_i^*}{e(1-\theta_i^*)}, \tag{11}$$

and reputations is

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

Proof. Consider group i. Combining equations (9) and (10), we obtain equations (11) and (12). Notice that in (11) the left-hand side is increasing in θ_i^* while the right-hand side is decreasing in θ_i^*, so that there exists a unique solution, possibly a corner solution with zero quality when $F_i(e) = 0$. ■

Note that the equilibrium is unique (up to mixing at the thresholds, which bears no consequence given the continuum of producers). Perhaps the main take-away is that there is *free-riding in all groups* of the Collective Reputation structure. This implies that group design cannot attain efficiency unless information is perfect, and will always exhibit free-riding.

Applying this proposition to Full Collective Reputation, with one single group, follows immediately.

Corollary 1 (**Equilibrium under Full Collective Reputation**)
Under Full Collective Reputation, there exists a unique equilibrium, such that a producer with $\theta \leq \theta^*$ chooses $q = 1$, and $q = 0$ otherwise. Average quality is $F(\theta^*) = \frac{e-\theta^*}{e(1-\theta^*)}$ and the reputation for unidentified quality is $\mu^* = (1-\theta^*)F(\theta^*)$.

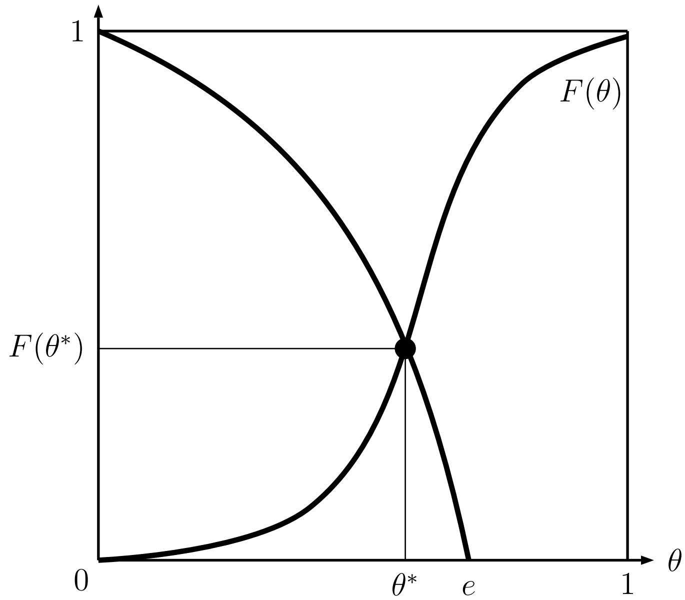

Figure 3: Equilibrium under Full Collective Reputation.

An illustration is provided in Figure 3. The downward-sloping curve depicts the right-hand side of (11), $\frac{e-\theta}{e(1-\theta)}$, which exactly coincides with the mixed equilibrium under IR. Indeed, by the indifference property, $q^*(\theta)$ corresponds to the average quality produced, such that, when correctly anticipated by the market, 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 producer with lower cost than θ produce high quality.

Welfare in equilibrium under FCR is easily computed as:

$$\begin{aligned} W_{FCR}(e) &= \int_0^{\theta^*} (1-\theta) f(\theta) d\theta \\ &= (1-\theta^*)F(\theta^*) + \int_0^{\theta^*} F(\theta) d\theta \\ &= \mu^* + \int_0^{\theta^*} F(\theta) d\theta. \end{aligned} \tag{13}$$

To understand in details the incentive effects of collective reputation the following comparative statics results are helpful.

Proposition 3 **(Comparative Statics under Collective Reputation)**
Consider group i with cost distribution $F_i(\theta)$. The comparative statics in e is:

- *more information increases incentives and average quality: θ_i^* and $F_i(\theta_i^*)$ increase in e.*

- *more information has ambiguous effects on reputation:*

  - *μ_i^* is increasing for e close to 0, and decreasing for e close to 1.*

  - *μ_i^* is single-peaked in e if F_i is log-concave: reputation is first-increasing, then decreasing when more information becomes available.*

Consider first-order stochastic decreases of F_i (i.e. stochastically lower costs):

- *quality increases when costs decrease: $F_i(\theta_i^*)$ increases when F_i decreases: .*

- *reputation increases when costs decrease: μ_i^* increases when F_i decreases.*

- *incentives in equilibrium decrease when costs decrease: θ_i^* decreases when F_i decreases.*

Proof. We suppress the index i for notational clarity. The first bullet is obtained by differentiating (11) with respect to e, which yields: $\frac{d\theta^*}{de} = \frac{1-(1-\theta^*)F(\theta^*)}{1-eF(\theta^*)+e(1-\theta^*)f(\theta^*)} > 0$ for all $e \in (0,1)$. For the second bullet, differentiating (12) 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 the third bullet, 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.

The fourth and sixth bullets come directly from inspection of (11) and (12). The fifth bullet comes from combining (10) with the fourth bullet. ∎

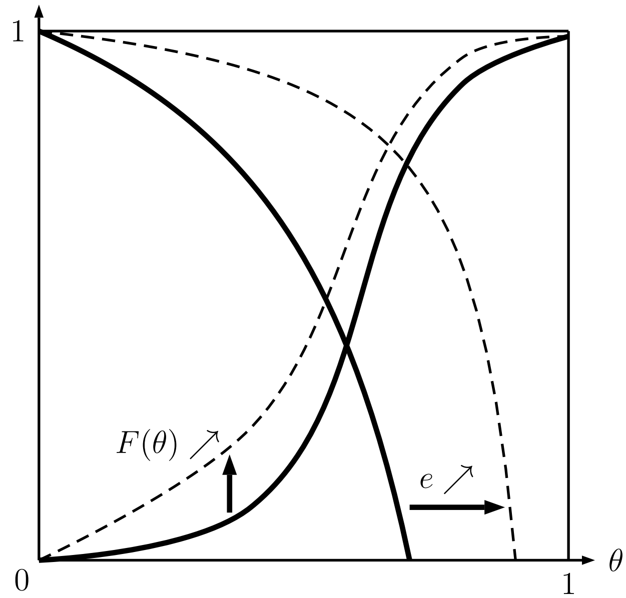

Figure 4: Comparative statics under Full Collective Reputation.

Incentives under Collective Reputation

From the above analysis, we can take a closer look at the implications of collective reputation on incentives within groups. First, observe that if a group produces quality, high cost producers of the group (above θ_i^*) always free-ride on the quality supplied by low cost producers in the group–a group cannot reach a reputation of 1 in equilibrium. Notice that every producer would like to expel the (other) high cost producers, since they all want a higher reputation μ_i^*. However, expelling the high cost producers would in turn lower incentives in the remaining group by increasing its reputation. Moreover, if low cost producers (below θ_i^*) would be separated from the group, incentives of the others would increase. We will study these aspects in the next section when analyzing admission thresholds.

From Proposition 3 we also see that all types of producers like some information to the extent that it increases reputation. Those who choose $q = 1$ because this increases their chances of recognition, those who chose $q = 0$ because it increases μ_i^* up to some point. For instance, if F is log-concave, producers unanimously agree to more information until the reputation reaches its peak.

In the following, we will compare incentives under individual and collective reputations.

3 Collective Reputation and Efficiency

In light of the previous comparative statics exercise, collective reputation has interesting incentives effects: excluding high cost members increases reputation, but reduces incentives. Conversely, excluding low cost members reduces reputation, but increases incentives. The effect of reshuffling across groups is far from trivial, and implications for overall quality and welfare in different collective reputation structures are not straightforward. To highlight the trade-offs inherent in different collective reputation structures, we will first start by comparing the two extreme structures, full collective reputation and individual reputations.

In this section, we will at times make use of the following assumption:

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

This assumption is used for technical convenience, and we point out where it plays a role when relevant. One part does however bear some important economic content. A positive f implies that the group contains the equivalent of commitment types, the zero-cost agents, who always produce quality. This will prove important, as it has been recognized before in other reputation models.

In this section, in order to economize on notations, we restrict attention the case of Full Collective Reputation. Adapting the results to any general Collective Reputation structures as defined previously is straightforward, with only minor qualifiers.

3.1 Full Collective Reputation versus Individual Reputation

Under FCR, the group's reputation induce high cost producers to free-ride. In contrast, under individual reputation, a producer cannot commit to providing very high quality, in the sense that a corresponding market belief would directly kill the incentive to do so: reputation-milking partially destroys incentives.

In particular, under FCR producers of type $\theta > \theta^*$ choose low quality with probability one, while under IR types θ with $\theta^* < \theta < e$ choose $q > 0$. However, Under FCR, low cost producers ($\theta \leq \theta^*$) provide high quality with probability one, while under IR, all producers (except 0 cost types) choose quality with probability $q < 1$. In terms of cost efficiency, this means that for the same given level of quality, cost under IR are higher than under CR. An immediate consequence is that IR can be more efficient than CR only if it means higher average quality.

For a given profile of individual qualities $q(\theta)$, let us define agregate quality as:

$$Q = \int_0^1 q(\theta)f(\theta)d\theta. \tag{14}$$

There are several aspects deserving attention. Quality produced in one or the other Reputation arrangement, the relative extent of free-riding and reputation milking, and in the aggregate, efficiency as measured by welfare. Of course these dimensions are all related. Comparison of quality depends finely on the properties of the type distribution. In order to get some preliminary insights, we study in details the case of a uniform costs distribution. Figure 5 shows equilibrium qualities in both structures and highlights the quality and cost trade-off between both structures: More quality from more efficient types under FCR, versus quality from a larger set of types (with more inefficient types) under IR.

With the uniform distribution, one can obtain close form solutions (derived in appendix) for total quantity produced.

$$Q^*_{IR}(e) = \theta^*(e) = \frac{1}{2} + \frac{1 - \sqrt{(1-e)(1+3e)}}{2e} \tag{15}$$

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

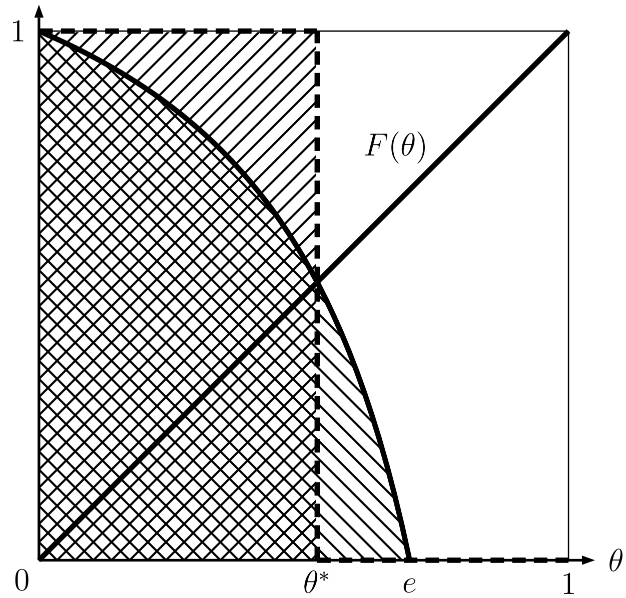

Figure 5: Quality provision: FCR vs IR (uniform distribution).

where we used (11) to obtain that $e - \theta^* = e\theta^*(1 - \theta^*)$. Hence we have just proved the following proposition.

Proposition 4 (**FCR vs IR with a uniform distribution**)
Suppose that F is the uniform distribution on the unit interval. Then Full Collective Reputation yields higher welfare than Individual Reputation.

In the case of the uniform distribution, Full Collective Reputations performs better than Individual Reputation for any experts' attention. However, proposition 4 does not generalize to arbitrary distributions.

The welfare comparison between collective and individual reputation generally depends non-trivially on the distribution of cost types. We can however say more about the ranking when information is poor. The next proposition asserts that when e is sufficiently small, FCR is better than IR in terms of welfare for any distribution.

Proposition 5 (**FCR vs IR with poor information**)
Suppose that F is regular. Full Collective Reputation strictly dominates Individual Reputa-

tion if e *is small enough. Formally,*

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

Proof. The first part is obvious: $W_{FCR}(0) = W_{IR}(0) = 0$. For the collective reputation case, using (13), direct computation shows:

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

Moreover, totally differentiating the equilibrium relation (11) with respect to e, one obtains:

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

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

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

For the Individual Reputation case, we obtain using (8):

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

Now, let

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

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

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

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

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

which allows to conclude since $f(0) > 0$ when F is regular. ∎

The takeaway from the proposition is that when quality is very rarely discovered by experts, Collective Reputation can help producers commit to higher quality. Despite the fact that the presence of zero-cost types creates some free-riding in the group, (Full) Collective Reputation is somewhat surprisingly strictly better than Individual Reputation when information is sufficiently poor. On way of understanding the results is that quality around $e = 0$ is of order 1 under CR, while it is only of order 2 under IR. The mixed strategy under IR falls very quickly in θ when e is small (indeed, $\lim_{e\to 0} \frac{\partial q^*(\theta)}{\partial \theta} = \lim_{e\to 0} -\frac{1-e}{e}$), while under FCR the equilibrium cutoff increases linearly with e around 0 (indeed, $\lim_{e\to 0} \frac{d\theta^*}{de} = 1$, so that the mass of effort is of order 1 for small e).

We can perform the same kind of exercise at the other extreme of information quality.

Proposition 6 (**FCR vs IR with almost perfect information**)
Suppose that F is regular. Individual Reputation strictly dominates Collective Reputation if e is high enough and the average cost is below 1/2. Formally,

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

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

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

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 $\lim\limits_{e\to 1} \frac{d\theta^*}{de} = +\infty$, h does not admit a Taylor expansion around $e = 1$, 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 welfare under individual reputation:

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

so that:

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

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. ■

The proposition hence shows the superiority of Individual Reputation when information is good enough and the costs are low enough. Intuitively, the fact that the welfare under IR increases around $e = 1$ proportionately to average costs underlines that better information means more effort in the whole population of producers. On the other hand, under FCR, only marginal producers at the threshold matter. Even though their additional value on welfare is small, if average costs is high, this means there are many of them.

For the gains associated spread over the whole population of producer to be higher than the gains at the top of the costs distribution, the average cost must be small enough, so that the gain for *all* types of producers associated with IR are more important

3.2 Admission thresholds

A common practice in constituting group is to set an admission standard. Schools run tests for allocating a limited number of seats and collective organizations of wine producers regulate entry. An immediate modeling of an admission threshold is to let the population of producers be split into two groups according to their cost parameter. We will consider an admission threshold σ, corresponding to admittance in the top group. The two groups have corresponding distributions F_1 and F_2, where group 1 corresponds to the set of producers with $\theta < \sigma$ (as a convention, producers with $\theta = \sigma$ belong to group 2).[15] These truncated distributions are thus:

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

Equipped with these cumulative distributions, one can directly apply the equilibrium characterization results obtained in Section 2.3.

Lemma 1 (**Cost distributions with an admission threshold**)
For any σ, F_2 first-order stochastically dominates F which first-order stochastically dominates F_1.
For $i = 1, 2$, 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: an increase in σ stochastically increase costs in both groups.

Proof. The first part is obvious given the formulas of truncated distributions. The second part follows simply from differentiating these formulas with respect to σ. ∎

Though rather intuitive, this lemma has several important implications when combined with the equilibrium comparative statics obtained in Proposition 3. Average qualities are

[15] In all that follows, some integrals are hence improper, since we consider an open set of producers for group 1. Since the highest cost producers in a group never exert effort, this is harmless.

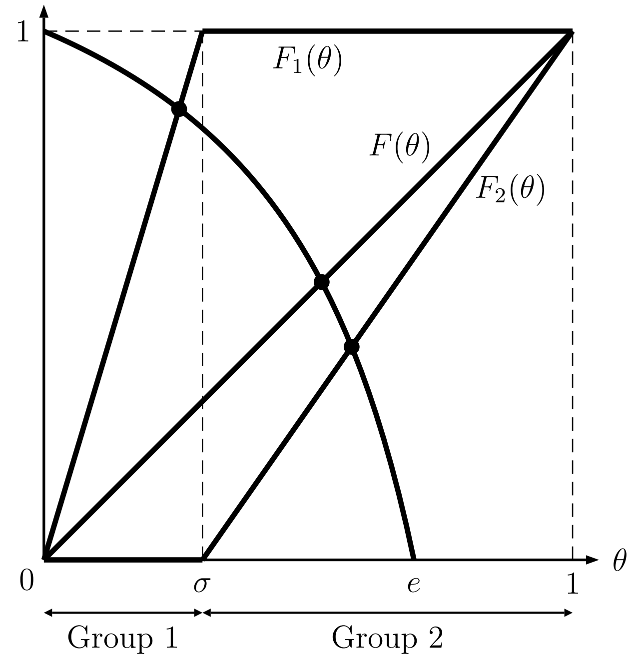

Figure 6: Equilibrium with an Admission Threshold (uniform distribution).

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 in group 1 compared to the case of a single group , incentives are reduced from equation (9). Hence, splitting the producers increases incentives for high cost producers at the costs of decreasing those for low cost producers. Thus, it is not straightforward whether one or the other solution is associated with more aggregate quality. Moreover, allocation of effort across producers with two groups is not optimal: between θ_1^* and σ, there is no effort, while producers between σ and θ_2^* again exert effort. Yet, the next result shows that a well chosen admission threshold improves efficiency.

Proposition 7 (**Admission Threshold**)
Suppose that F is regular. There exists an admission threshold $\sigma > 0$ such that the welfare with two groups separated at σ is higher than the welfare with Full Collective Reputation.

Proof. We abuse slightly notation by defining the welfare with cutoff σ as the function:

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

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 1 to the truncated distributions, we obtain by differentiating (11) with respect to σ in both groups:

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

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

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

such that the second term in the welfare derivative is positive. Moreover, since when $\sigma = 0$ we have $\theta_1^* = 0$, $\frac{d\theta_1^*}{d\sigma}|_{\sigma=0} = \frac{ef(0)}{ef(0)} = 1$ so that the first term in the welfare is zero. In total, we thus have $\frac{dW}{d\sigma}|_{\sigma=0} > 0$, so that there exists an interior cutoff σ strictly improving on the single group. ■

Thus, with any regular distribution,[16] FCR can be improved on by introducing an admission threshold into an elite group. Note however that a cutoff σ with $e \leq \sigma < 1$ always yields lower welfare than FCR: For $\sigma = e$, there is no effort in group 2, and since $\theta_1^* < \theta^*$, $F(\theta_1^*) < F(\theta^*)$, so that quality is lower with a cutoff $\sigma = e$ than under full reputation. Since it is efficient for all θ with $\theta_1^* \leq \theta \leq \theta^*$ to exert effort, welfare with a cutoff $\sigma = e$ is lower

[16] We only use the fact that $f(0) > 0$ in the proof of Proposition 7.

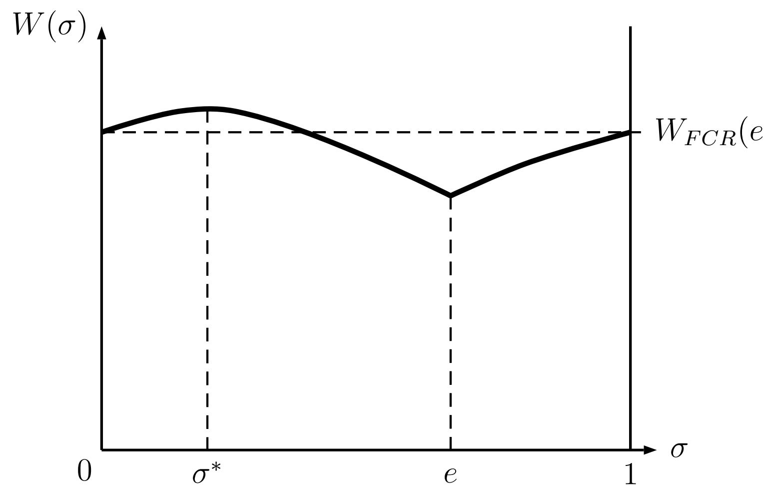

Figure 7: Welfare with an Admission Threshold (uniform distribution).

than welfare under full reputation. Now note that welfare is increasing in σ for $\sigma \geq e$ for the following reasoning: θ_1^* is increasing in σ, and there is no change in group 2. Hence, average quality is increasing in σ for $\sigma \geq e$ with only the most efficient producers increasing effort such that welfare increases.

The use of admission thresholds should be an informed choice: a threshold can actually decrease welfare if unwisely chosen. Figure 7 illustrates that in the case of the uniform distribution.

The last proposition shows that going from Full Collective Reputation to an organization of producers that splits producers into two groups by a cutoff increases welfare. The logic behind it is that (1) the free-riding in the high-cost group is reduced such that quality in this group is increased by more than it is reduced in the other group (higher free-riding among low-cost producers), and (2) the potential allocation inefficiency from more inefficient types exerting effort is not large.

Interestingly, in light of the previous proposition, merging of groups can also be an efficient solution. If the population is split with an outdated threshold, perhaps calibrated on a distribution that is not accurate anymore, merging the groups could be undertaken instead of attempting to change a long standing rigid admission scheme.

3.3 Free Group Entry

[Just explain (maybe with a proposition) that when free-riders are added to the group, without size limit, i.e. for instance producers with costs higher than 1, then the reputation converges to zero, and the only hope is to be detected by the market. This creates maximal incentives by reducing reputation]

Hence the equilibrium converges to $\theta^* = e$, which obviously strictly dominates IR.

4 Sustainability of Collective Reputation

4.1 Collective Reputation Unravelling

In the previous section, we have shown that 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 individual producers would prefer to stay in a group rather than work under their individual reputations. In fact, individual incentives are such that the lowest cost producers in any group producing quality would always prefer to split from the group such that collective reputation would unravel, which we show in the next proposition.

Proposition 8 (**Group unraveling**)
In a collective reputation structure, any group producing quality unravels to individual reputation.

Proof. 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 $U_{ind}(\underline{\theta}_i) = \mu(\underline{\theta}_i) = 1 - \frac{\underline{\theta}_i}{e}$ under individual reputation and $U_{coll}(\underline{\theta}_i) = \mu_i^* + \theta_i^* - \underline{\theta}_i = 1 - \frac{\theta_i^*}{e} + \theta_i^* - \underline{\theta}_i$ under collective reputation. We have $U_{ind}(\underline{\theta}_i) - U_{coll}(\underline{\theta}_i) = (\theta_i^* - \underline{\theta}_i)(\frac{1}{e} - 1) > 0$. ■

The problem underlying the result–that lowest cost producers in any group producing quality always have a higher payoff under individual reputations–is illustrated in Figure 6 for

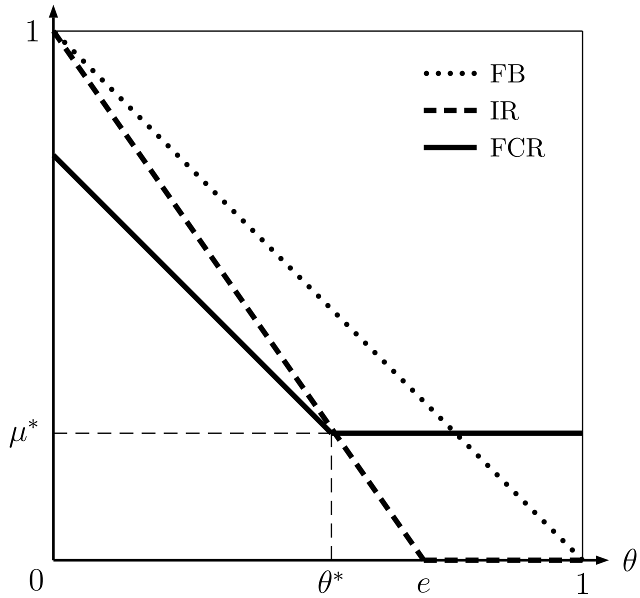

Figure 8: Individual payoff: FB, IR and FCR.

FCR with the uniform distribution.[17] Note that, contrary to unraveling with pure Adverse Selection, which is efficient, unraveling here with Moral Hazard can be inefficient.

Thus, the question is whether a group can implement an internal incentive system to sustain collective reputation in order to reap its benefits. Clearly, when types are observed within a group and transfers between group members can be type-dependent, 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 types are not observed within the group and transfers cannot be type-dependent, collective reputation can be sustained. Even more, a simple non-type-dependent transfer scheme can ensure participation and increase incentives of group members to the First Best quality level.

[17] This is precisely a case in which FCR yields higher welfare (and producer surplus) than IR, as stated in Proposition 4, so unraveling is suboptimal.

4.2 Efficient Group Incentive Schemes

We consider group incentive schemes set in place by a Collective Organization that can pass contracts with its members. The only information available for contracting is the market price received.[18] Hence an incentive scheme comprises (at most) two instruments based on market price. We will consider the following group incentive scheme, without loss of generality: a participation fee t is levied on each producer and a bonus for identified high quality s. The bonus s increases incentives in the group on top of price incentives. Of course, the equilibrium depends on the subsidy level s, but note that t does not influence the equilibrium once participation is ensured. Hence let θ_s and μ_s denote the threshold and reputation under a subsidy s. The corresponding incentive constraint now writes:

$$e(1+s) + (1-e)\mu_s - \theta - t \geq \mu_s - t \tag{16}$$

while Bayes' rule applies to $F(\theta_s)$. Combining these two relations yields the following new equilibrium relation taking s into account:

$$F(\theta_s^*) = \frac{e(1+s) - \theta_s^*}{e(1 + es - \theta_s^*)} \tag{17}$$

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 individual reputation payoff:

$$(1-e)\mu_s^* - \theta - t + e(1+s) \geq \mu^*(\theta) \tag{18}$$

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

$$t \geq es. \tag{19}$$

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

[18] When types are not observed in a group, one can easily see that the only mechanisms available have to be pooling, i.e. it is impossible to tailor incentives to the type of a producer when it is not observable.

Proposition 9 (**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* $s^* = \frac{1}{e}$ *for an identified high quality product.*

Proof. Note that as previously, equation (17) has a unique solution θ_s^*. It is immediate to check that $\theta_s^* = 1$ when $s = \frac{1}{e}$. As a consequence, $\mu_s^* = 1$. A producer's profit is then $1 - \theta$. The participation constraint (18) can now be written as $\mu^*(\theta) + \theta \leq 1$, or by substituting from proposition 1: $1 - \theta\left(\frac{1}{e} - 1\right) \leq 1$, which always holds. Note finally that this scheme respects budget-balance, since the fee of 1 is levied on all producers, while the subsidy is paid with average probability e, for a unit cost of $\frac{1}{e}$. ■

The idea of the scheme is to grant producers the right to sell under the collective reputation, and to entitle them to bonuses paid by the collective organization. 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. With this "bonus club" scheme, producers all receive in expectation their first-best payoff of $1 - \theta$, which they can not 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 collective organization in charge of marketing. The collective organization can then redistribute with distorted prices, to calibrate incentives. Not only the collective reputation price μ can be kept by the collective organization to increase incentives, it can in addition be redistributed to high quality producers to increase incentives even further.

Proposition 10 (**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 buying a lottery ticket that is worthless without effort, but worth 0 with probability $(1 - e)$ and $1/e$ with probability e with effort. 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.

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 in Bordeaux–admittedly a collectively successful region. Likewise, famous alumni are more often praised and rewarded by their alma mater than average students.[19]

[19] Financially successful alumni sometimes even repay the reputation they have acquired through donations–this might constitute another way of benefiting from exposition from their alma mater.

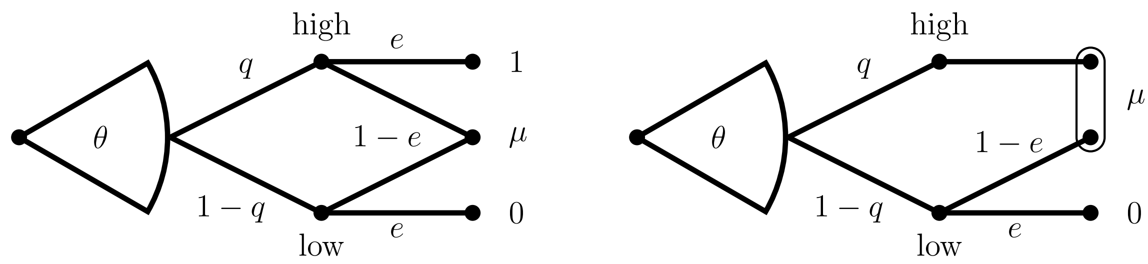

Figure 9: Unbiased (left) and hostile (right) information structures.

5 Discussion and extensions

One assumption we have made so far deserves a substantial discussion. While we have studied extensively friendly environments, several results do not carry over to information structures that are neutral or hostile. We sketch in this section the main differences and present the counterparts to our main results.

5.1 Selective vs non-selective bias in information revelation

So far our analysis has relied on a particular form of selective bias in quality revelation: only good quality can be made public with some probability. As argued already, this assumption fits well numerous examples, such as wine or education. In other circumstances, the nature of feedback is instead hostile: only bad news tend to reach the public. For instance, various scandals such as contaminated beef or the Volkswagen diesel emission fraud corresponds to cases where in business as usual, no news is produced, and only in some instances of misbehavior or cheating on quality does the news reach the public. Before considering this alternative information structure, we however would like to point out that selective news transmission is necessary for collective reputation to play a role in efficiency. Consider an alternative model in which monitoring is all-or-nothing, so that quality is either revealed with probability e, or not revealed, as represented on Figure 9. Then, for any μ, the incentive constraint now writes as follows:

$$e + (1-e)\mu - \theta \geq e.0 + (1-e)\mu, \tag{20}$$

which simply amounts to

$$\theta \leq e.$$

Hence each producer's incentives are independent of the collective reputation design.[20] It does not mean of course that a producer is indifferent to his group membership: μ does enter his payoffs, and a producer prefer to belong to groups with higher reputation. But in terms of efficiency, group composition (and hence group design) is irrelevant with an unbiased information structure.

5.2 Collective Reputation in a Hostile Environment

The seminal model by Tirole (1996) features a particular news selection bias: only bad news (corruption, in his main example) can surface–the environment is (purely) hostile. As already discussed, this is a key assumption in generating multiple equilibria and explaining persistence of bad equilibria in dynamic models. This insight is the core of Tirole's analysis, which we will not repeat it here. On the contrary, we would like to focus on other implications of hostile environments for collective reputation design, and contrast them with the friendly case.

The key difference is that now incentives come from the 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:

$$\mu - \theta \geq e.0 + (1-e)\mu, \tag{21}$$

which amounts to

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

The crucial implication is that reputation and incentives now go hand in hand. A higher

[20]This includes the case of individual reputation, since the incentive constraint is identical. In that case as well all agent with $\theta \leq e$ choose $q = 1$.

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 the market believes 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 correct belief $\mu(\theta)$ of the market. 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:

$$F(\theta^*) = \frac{(1-e)\theta^*}{e(1-\theta^*)}. \tag{22}$$

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 (22) 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 all 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 $\sigma = e$.[21] By applying Lemma 1 to our new equilibrium relation, we obtain that there always exists an equilibrium in which 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 they could reach under IR.[22] The rationale for the admission threshold here is however different from that in the friendly case. While in a friendly case 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

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

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

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.

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[23] 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.

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

As a positive final note, it is important to bear in mind that a grain of friendliness, that is 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.

5.3 Endogenous attention of experts

Another important aspect that we have kept exogenous in the model is the intensity of attention granted by experts, as captured by the parameter e. If we allow attention to be a function of group costs, additional interesting phenomena arise. Suppose for instance that experts like slightly more to discover high quality–say because this increases their audience, which is the source of their revenue. If experts seek good news, and must decide first which group to inspect, then they will prefer to inspect groups with higher expected quality (hence probably lower costs groups). Producers in this group are in turn incited to supply more quality under more intense scrutiny. In other words, experts' attention creates a self-fulfilling phenomenon. In addition, if several experts choose sequentially which group to inspect, they will more likely inspect groups who are already under more scrutiny, because this means higher quality in equilibrium. As a result, endogenous experts' attention create both self-fulfilling prophecies regarding quality, and herding of experts in group choice. Examples in the wine industry abound, where top quality regions attract disproportionately more attention than the others, and in education and employment, where recruiters rely on school's name to select applicants to interview. Admission thresholds under endogenous attention magnify elitism and tiering, and add to the reputation gap. But even two identical population of producers can end up in very different equilibria given this social coordination mechanism and self-fulfilling prophecies associated with attention.[24]

[24] Note that this argument goes further than the usual arguments of multiple equilibria in the statistical discrimination literature originated by Arrow (1973), where the information structure is fixed, as in Coate and Loury (1993) for instance. On endogenous attention and the nature of feedback, see also Halac and Prat (2016).

6 Conclusion

Collective reputations are pervasive in many industries, and often they help market players to solve informational problems that plague exchanges of quality products. In this paper we present a new model of collective reputation and focus on positive feedback on quality, i.e. where producers are facing a friendly informational environment. We analyze the incentive effects of collective reputation and show that free riding under collective reputation may be less detrimental to quality and welfare than reputation-milking under individual reputations. Our analysis highlights two important roles of Collective Reputations versus Individual Reputations: First, a reshuffling of quality incentives that can entail more cost-efficient production. Second, the possibility to provide group incentives on top of market incentives—without any additional information requirements—to boost quality provision.

References

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

Arrow, K. J. (1973). The theory of discrimination. In O. Ashenfelter and A. Rees (Eds.), *Discrimination in Labor Markets*, pp. 3–33. Princeton, N.J.: Princeton University Press.

Baldwin, R. and J. Harrigan (2011). Zeros, quality and space: Trade theory and trade evidence. *American Economic Journal: Microeconomics 3 (2)*, 60–88.

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. *The 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.

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

Crespi, J. and S. Marette (2002). Generic advertising and product differentiation. *American Journal of Agricultural Economics 84(3)*, 691–701.

Crozet, M., K. Head, and T. Mayer (2012). Quality sorting and trade: Firm-level evidence for French wine. *Review of Economic Studies 79*, 609–644.

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. Finkelshtain, A. Simhon, and N. Yacouel (2018). Collective brands. *International Journal of Industrial Organization, forthcoming*.

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.

Halac, M. and A. Prat (2016). Managerial attention and worker performance. *American Economic Review 106*(10), 3104–32.

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. *The B.E. Journal of Theoretical Economics - Contributions 9*, art. 27.

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

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

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 (2016). The benefits of collective reputation. Discussion paper no. 2068, Cowles Foundation - Yale University.

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

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.

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