Discussion Paper Series – CRC TR 224

Discussion Paper No. 281
Project B 03

Collective Brand Reputation

Volker Nocke [1]
Roland Strausz [2]

March 2021

[1] University of Mannheim, University of Surrey, and CEPR. Email: volker.nocke@gmail.com
[2] Humboldt Universität zu Berlin and CEPR. Email: roland.strausz@hu-berlin.de

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Collective Brand Reputation*

Volker Nocke† Roland Strausz‡

March 23, 2021

Abstract

We develop a theory of collective brand reputation for markets in which product quality is jointly determined by local and global players. In a repeated game of imperfect public monitoring, we model collective branding as a pooling of quality signals generated in different markets. Such pooling yields a beneficial informativeness effect for the actions of a global player present in all markets, but also harmful free-riding by local, market-specific players. The resulting tradeoff yields a theory of optimal brand size and revenue sharing, applying to platform markets, franchising, licensing, umbrella branding, and firms with team production.

Keywords: Collective branding, reputation, free riding, repeated games, imperfect monitoring.

*We thank Heski Bar-Issac, Helmut Bester, Jay Pil Choi, Paul Heidhues, George Mailath, Moritz Meyer-ter-Vehn, Benny Moldovanu, Aniko Öry, Martin Peitz, and Juuso Välimäki for helpful comments. We also thank seminar audiences at Edinburgh, Paris-Dauphine, NUS, TSE, and the CEPR Virtual IO Seminar as well as participants at the 1st MaCCI-EPoS-CREST Industrial Organization Workshop (Mannheim), the 2019 Meeting of the *Industrieökonomischer Ausschuss* (Bern), the 5th Workshop on Relational Contracts (Madrid), the 12th Workshop on the Economics of Advertising and Marketing (Porto), the 1st Cornell University Applied Theory Conference, the 3rd University of Bergamo Workshop on Advances in Industrial Organization, the 2019 Asia-Pacific Industrial Organization Conference (Tokyo), and the 2020 North American Winter Meeting of the Econometric Society (San Diego). We gratefully acknowledge financial support from the German Research Foundation (DFG) through CRC TR 224 (Project B03) [Nocke] and CRC TR 190 (Project B02) [Strausz].

†University of Mannheim, University of Surrey, and CEPR. Email: volker.nocke@gmail.com.

‡Humboldt Universität zu Berlin and CEPR. Email: roland.strausz@hu-berlin.de.

1 Introduction

Most products are experience goods in that product quality is not observable prior to consumption. When making their purchasing decisions, consumers therefore take into account the long-lived reputation of a brand, formed through past consumption experiences, word of mouth, internet reviews, and the like, and facilitated by trade-marks and logos. Managing its brand reputation is therefore crucial to a firm's success.[1] Controlling quality at the brand level is, however, challenging because it is a coarse measure that, in general, depends on the collective decisions of different types of agents. For instance, platform brands such as Amazon, Ebay, and Uber, or franchise brands such as McDonald's and Burger King have some control over consumers' experience but, to a large extent, this experience also depends on the individual sellers, drivers, and outlets, who ultimately service the consumer locally.[2] Similarly, the quality of manufacturing goods such as cars or beers depends not only on the decisions taken by the headquarter but also on those by local plant managers. A crucial feature of brand reputation is, therefore, that it is a collective outcome of actions taken not only by "global" players, who impact the quality of the entire line of products, but also "local" players, each of whom affects quality of only a subset of products.

Collective branding effectively involves pooling the quality signals from the entire product line. While this may make it easier to provide incentives to global players, it comes at the cost of inducing free-riding behavior by local players. This raises a number of important questions. In particular, under what conditions is a collective brand reputation optimal and when would it be better to sell the different products under different brand names? For example, is it better to sell six different products under a single, two, three, or six independent brand names? More generally, what is the optimal size of the collective brand? Moreover, what instruments can be used to mitigate free-riding incentives and how large are the remaining inefficiencies?

To address these and related questions, we develop a framework to analyze collective brand reputation. More specifically, we analyze an infinitely repeated hidden action game of imperfect public monitoring. There are n markets. In each market, two players jointly produce a good over an infinite number of periods. One of these players – the local player – is active in only one market, whereas the other player – the global player – is active in all n markets. The good is of high quality in a given market only if both players in that market

[1]Brand Finance, which produces annual rankings of companies based on brand/intangible value, reports that, on average, the total intangible value of a global top ten company accounts for over 85 percent of such a firm's market value (https://www.visualcapitalist.com/intangible-assets-driver-company-value/).

[2]Nosko and Tadelis (2015) document such problems in platform markets. Blair and Lafontaine (2005) highlight the franchisor's problem of maintaining consistent quality across franchisees, and conclude that the empirical evidence "seems consistent with the concept that free riding, or individual profit maximization by opportunistic franchisees, is an issue in franchised chains" (p.137).

exert effort. Neither effort nor quality are observable but, at the beginning of each period and for each market, a noisy (binary) signal of last period's quality is realized. Throughout, we focus on (symmetric) Perfect Public Equilibria (PPE).

The n goods can either be sold under separate brand names or under a single, collective brand name. Brand reputation is modeled according to a key tenet of marketing: consumers identify quality with the reputation of the good's brand.[3] Under independent branding (when the products in the various markets are sold under different brand names), behavior in each market therefore depends only on the signal realizations in that market. By contrast, under collective branding (when the products are all sold under the same brand name), the signal realizations in the different markets all pertain to the same brand. As a result, behavior in each market depends on the pooled quality signals from all markets rather than the individual signals.

Under independent branding, we show that high quality provision can be sustained if the discount factor is sufficiently large. In the best PPE, equilibrium play stochastically transits to the worst PPE (where no effort is provided) after a bad signal realization. Optimal revenue sharing takes the simple form of proportional rewards: each player's optimal revenue share is equal to the share of the total effort cost that needs to be borne by that player. The resulting outcome is the same as if one player had to make both effort decisions, bearing all of the costs in return for all of the revenue. However, because monitoring is imperfect (and consumers are short-lived), the Folk Theorem does not obtain: the on-path breakdown probability is bounded away from zero, no matter how large the discount factor.

Whether, compared to independent branding, collective branding reduces or increases the inefficiency, depends on our model's novel trade-off between the *informativeness effect* of pooling signals across markets and the *free-rider effect* of making continuation values (of local players) depend on signals that they cannot affect. It is instructive to isolate these effects by first considering two polar cases.

In the polar case in which the global player bears all effort costs (so that local players can be incentivized for free), collective branding is unambiguously beneficial. We obtain this result by first showing that the outcome under independent branding can be replicated with an appropriate choice of transition probabilities to the worst PPE, with that probability being strictly increasing in the number of good signals. We then show that this outcome can be improved upon, because the best PPE exhibits the following cut-off structure: There exists a cutoff $\tilde{s}$ such that the transition occurs for sure if the number of good signals is below $\tilde{s}$, and the transition does not occur at all if the number of good signals exceeds $\tilde{s}$.

The optimality of the cut-off structure results from the beneficial informativeness effect of

[3] See, for example, the textbook by Kotler (2003, p.420): "A brand is essentially a marketer's promise to deliver a specific set of features, benefits, and services consistently to the buyers."

a collective reputation: Even though signals are independent across markets and the global player can choose in how many markets to shirk, pooling signals and punishing only if a sufficiently large number of them is bad, reduces the on-path breakdown probability, thereby decreasing the inefficiency associated with imperfect monitoring. In this polar case, the equilibrium value is therefore increasing in the collective brand size n: it is optimal to sell all products under a single brand rather than multiple ones. Moreover, as long as the discount factor is above the critical level for sustaining high effort under perfect monitoring, efficiency obtains in the limit as n becomes large.

In the other polar case in which the local players bear all effort costs (so that the global player can be incentivized for free), collective branding is unambiguously harmful. This is due to the free-rider effect which arises because, under collective branding, each local player's continuation value depends on signal realizations from other markets that he cannot influence. We show that, in this case, the outcome under independent branding can, at best, be replicated, and that this replication is only possible if the on-path punishment probability is small, requiring a sufficiently large discount factor.

Turning to the generic case in which each type of player bears some of the effort costs, proportional rewards (which are optimal under independent branding and in both polar cases under collective branding) fail to mitigate the free-rider effect: under proportional rewards, the best outcome in terms of joint value is in fact identical to that in the polar case in which the local players bear all effort costs.

We subsequently show that, in the generic case, optimal revenue sharing involves giving a more-than-proportionate share of the revenue to the local players, thereby mitigating the free-rider effect. For sufficiently large discount factors, the best PPE under collective branding allows sustaining a lower on-path breakdown probability than the best PPE under independent branding. There is, however, a finite upper bound on the brand size that allows sustaining effort, and so the optimal brand size is finite. As the discount factor becomes large, this optimal finite brand size increases without bound. In the limit as both the discount factor and the brand size become large, the inefficiency from imperfect monitoring does, however, not vanish. The remaining inefficiency is equal to the share of the effort cost borne by the local players multiplied by the inefficiency under independent branding. Hence, with optimal revenue sharing, it is as if, in that limit, one achieves the best of both worlds: collective branding for the global player and independent branding for the local players.

Related literature. The unique feature of our paper is to study collective reputation in the presence of both global and local players.[4] The key tradeoff underlying our results is between the informativeness effect (which is beneficial in the presence of a global player) and

[4]See Bar-Isaac and Tadelis (2008) for a survey of the literature on seller reputation.

a free-rider effect (which is detrimental in the presence of local players). While the economics literature has studied variants of these two effects in isolation, our paper is the first to study their interactions and economic implications.

Studying cooperation in a repeated prisoner's dilemma game with imperfect monitoring, Matsushima (2001) was, to our knowledge, the first to identify the beneficial informativeness effect that underlies our paper. Cabral (2009) shows that this informativeness effect may render umbrella branding optimal, while Cai and Obara (2009) study it as a driver of horizontal integration. These papers effectively only consider global players, thereby abstracting from free riding on a collective reputation.

By contrast, Tirole (1996), Fishman et al. (2018), Neeman et al. (2019) study free riding problems associated with a collective reputation in settings with only local players.[5] These papers however do not consider our beneficial informativeness effect of a collective reputation, and therefore address different economic forces. Considering a repeated matching environment with overlapping generations, Tirole (1996) shows that a collective reputation may lead to a persistent stigmatization of new generations due to shirking by some earlier generation. In his framework, there are no inherent benefits from a collective reputation. In a two-period model with persistent investment and different types of firms, Fishman et al. (2018) study free-riding on a collective brand reputation, but with the benefit that the collective brand has the assumed ability to select its members based on their investment decisions and/or types. Neeman et al. (2019) point out that a collective reputation may serve as a commitment device. Trading off this effect against the free-rider effect, they therefore study a different trade-off from ours.

Focusing on brand management, our study is related to the literature on co-branding, brand extension, and umbrella branding (e.g., Kotler, 2003). With respect to this extensive literature, our contribution is to study potential free-riding problems which, in our view, are endemic to such settings.[6] While most work in this literature focuses on reputation models with hidden information (e.g., Wernerfelt (1988), Choi (1998), Cabral (2000), Miklos-Thal (2012), Moorthy (2012)), our work is more closely related to studies that analyze umbrella branding in a moral hazard framework. Building on Klein and Lefler (1981), Andersson (2002) shows that in a repeated game of moral hazard but perfect monitoring, a single brand name that pools the reputation across independent markets is helpful only if markets display asymmetries. Hakenes and Peitz (2008) and Cabral (2009) highlight that, with imperfect monitoring, pooling reputation can be beneficial even when markets are symmetric.

Moreover, our paper contributes to the literature on the management of moral hazard in

[5] Winfree and McCluskey (2005) and Fleckinger (2014) also study collective reputation, but in a framework in which consumers observe the product's (collective) quality at the time of purchase.

[6] See Castriota and Delmastro (2012) for an empirical study of the importance of collective reputation in wine markets.

team production, pioneered by Holmström (1982). In our setup, there are two types of team production: physical team production within a market and reputational team production across markets. Since the focus of our analysis is on the reputational team production problem, we assume that the market-specific effort choices of the local and global players are perfectly complementary so that these players can fully resolve their physical team production problem.

Our results also provide insights into the classical quality management problem in franchising and licensing. In franchising (licensing), we can view the franchisor (licensor) as a global player, while the outlets (licensees) are local players. Practitioners and legal scholars have pointed out the importance of free-riding problems in these contexts. For instance, Hadfield (p.949, 1990) notes that the individual "franchisee is inclined to make decisions about how much effort to put into the business based on the profits that will accrue directly to her in her own outlet," whereas "customers make judgments about the quality of the entire franchise system based on their experience at an outlet".[7] Similarly, Klein and Saft (p.349ff, 1985) remark that the "franchise arrangements create an incentive for franchisees to shirk on quality," further pointing out that "the individual franchisee directly benefits from the sales of the low quality product, and the other franchisees share in the losses caused by decreased future demand".[8] In the context of trademark licensing, Calboli (2007) explains that, legally, "trademarks are protected only as conveyors of information about the products which they identify and as symbols of commercial goodwill" (p.357) and points out the free-riding problem that licensees' "lack of direct ownership of the mark could make them less interested in the long-term success of the products" (p.360).

To our knowledge, our paper is the first to formally model, and rigorously analyze, this classical problem.[9] Our result that, without properly calibrating the revenue shares, a collective reputation destroys any benefits from pooling reputations confirms that due to free-riding "the value of the trademark will suffer dramatically" (Hadfield, 1980). Yet, our results also

[7]A concrete example is the Burger King scandal in Germany in 2014. After an undercover report exposed severe problems of poor hygiene in outlets in Cologne, Burger King tried to put the blame on the individual outlets but German consumers associated the negative report with the Burger King brand as a whole rather than its local franchisee in Cologne. Similarly, in Kentucky Fried Chicken Corp. v. Diversified Packaging Corp., 549 F.2d 368, 380 (1977), the Court observed "A customer dissatisfied with one Kentucky Fried outlet is unlikely to limit his or her adverse reaction to the particular outlet; instead, the adverse reaction will likely be directed to all Kentucky Fried stores. The quality of a franchisee's product thus undoubtedly affects Kentucky Fried's reputation and its future success."

[8]In the context of licensing, the court in Siegel v. Chicken Delight, 448 F.2d 43, n.38 (1971) observed that "the licensor owes an affirmative duty to the public to assure that in the hands of his licensees the trade-name continues to represent that which it purports to represent." Klein and Saft (p.349ff, 1985) interprets this view as expressing "a legal obligation for quality maintenance in a system involving many producers operating under a common trade name."

[9]Extensively discussing the free-riding problem in franchising, Blair and Lafontaine (2005) capture the bare essentials of this problem in a highly stylized model that abstracts from any reputational concerns.

show that a franchisor can partially mitigate free-riding problems and thereby maintain the trademark's value by shifting revenue streams from himself to the franchisees. This reinforces the insight of Bhattacharyya and Lafontaine (1995) that revenue sharing is crucial for controlling double moral hazard problems in franchising in that free-riding problems associated with a collective reputation drastically affect optimal revenue shares.

Plan of the paper. In the next section, we present the model. This is followed, in Section 3, by the equilibrium analysis of independent branding. In Section 4, we first study the polar case of collective branding in which the burden of effort is borne entirely by the global player. We then turn to the other polar case in which all of the burden of effort is borne by the local player. The analysis of these polar cases is instructive for analyzing the generic case in which local and global players share the effort cost. Section 5 addresses the comparative statics in the brand size n and obtains results concerning the maximum implementable brand size, $\bar{n}$, the optimal brand size, $\hat{n}$, and (in)efficiency results for limiting cases. We conclude in Section 6. We collect all proofs in Appendix A.[10]

2 The Model

We consider an infinitely repeated game of imperfect public monitoring in discrete time $t = 0, 1, \ldots$. There are $n \geq 2$ symmetric markets, indexed by $i = \{1, \ldots, n\}$ with one long-lived global player, G, and n long-lived local players, L_i. For each period t, production in a market i requires the market-specific binary input, $e^t_{G,i} \in \{0, 1\}$, of the global player G and the binary input, $e^t_{L,i} \in \{0, 1\}$, of the market-specific local player L_i. The good produced in market i is sold to a (representative) market-specific short-lived consumer C_i.

Production technology. The quality q^t_i of good $i \in \{1, \ldots, n\}$ in period $t \in \{0, 1, 2, ..\}$ is either high, 1, or low, 0, and depends on the contemporaneous effort choices of G and L_i. In particular, it is equal to one if and only if both G and L_i put in effort for that good (i.e., $e^t_{G,i} = e^t_{L,i} = 1$), and zero otherwise. The aggregate cost of effort for producing high quality in a specific market is $c > 0$, of which G incurs the share λ_G and L_i incurs the remaining share $\lambda_L = 1 - \lambda_G$. That is, λ_G represents the importance of the global player's effort cost relative to that of the local player, and the effort cost of G is $e^t_{G,i}\lambda_G c$ and that of L_i is $e^t_{L,i}\lambda_L c$. Markets are symmetric in that these shares do not differ across markets.

[10]By applying the abstract methods of decomposability and self-generation developed in Abreu, Pearce, and Stacchetti (1990), we study, in Appendix B, asymmetric PPE in the case of $n = 2$ markets, addressing the robustness of our results to asymmetric equilibrium outcomes.

Timing. The infinitely repeated game starts after the long term players set, for each good i, the revenue shares π_G and $\pi_L = 1 - \pi_G$ that accrue to G and L_i, respectively. At the beginning of each period $t \geq 1$, before effort choices are made, there is a binary signal $s_i^t \in \{0, 1\}$, providing noisy information about the good's quality in the previous period.[11] Without loss of generality, in period 0, the value of the signal is normalized to one: $s_i^0 \equiv 1$. As we formalize below, the collective branding decision determines the (public) observability of these signals. The realization of the market-specific signal s_i^t depends on the previous period as follows: if $q_i^{t-1} = 1$, then $s_i^t = 1$ with probability $1 - \alpha$; similarly, if $q_i^{t-1} = 0$, then $s_i^t = 0$ with probability $1 - \beta$. The parameters $\alpha \in (0, 1)$ and $\beta \in (0, 1 - \alpha)$, measure the noisiness of the signal and represent the probabilities of type II and type I errors, respectively. Before effort choices are made, there is also the realization of an independent public randomization device $r^t \in [0, 1]$, uniformly distributed over the unit interval, and independent over time.[12] After the effort choices have been made, the consumer in market i, C_i, decides whether to buy good i ($b_i^t = 1$) or not ($b_i^t = 0$).

Payoffs. The consumer's utility from consumption equals the good's quality level. As we focus on equilibria in which all long-lived players exert effort so that the consumer rationally expects high quality, we fix the price of the good to 1.[13] As agreed upon in an initial stage prior to the repeated game, the global player G and local player L_i receive shares π_G and $\pi_L = 1 - \pi_G$ of the revenue from selling good i. Assuming that signals s_i^t are non-contractible, these shares cannot condition on the signal realizations and, focusing on symmetric equilibria, they are uniform across markets. We discuss the feasibility of more elaborate ways of revenue sharing under different modeling assumptions in the conclusion.

A natural division of the revenue is to set a player's reward share equal to his cost share, $(\pi_G, \pi_L) = (\lambda_G, \lambda_L)$. We refer to this sharing rule as *proportional rewards.* For the case in which costly effort is needed from both the global and local players, $\lambda_G, \lambda_L > 0$, we define the *reward-to-cost-share ratio* of player j as $\gamma_j \equiv \pi_j / \lambda_j$. Proportional rewards then correspond to $\gamma_G = \gamma_L = 1$, while the identities $\lambda_L = 1 - \lambda_G$ and $\pi_L = 1 - \pi_G$ imply the accounting identity $\gamma_G = \pi_G / \lambda_G = (1 - \lambda_L \gamma_L)/(1 - \lambda_L)$.

Summarizing, the period-t profit of a long-lived player $k \in \{G, L_i\}$ in market i is equal to

$$b_i^t \pi_k - e_{k,i}^t \lambda_k c.$$

[11] For notational convenience, we define $s_i^0 = 1$.

[12] The public randomization device simplifies the exposition of our results; none of our results require the existence of such a signal. Footnote 16 makes this explicit.

[13] In our formal modelling of the repeated game, we treat this price as fully exogenous. Equivalently, we could have assumed—as is commonly done in the literature on umbrella branding—that the price is equal to consumers' willingness to pay for the good, e.g., because multiple (identical) consumers bid for the good in a (second-price) auction.

The long-lived players discount profits with factor $\delta \in (0,1)$. The payoff of the (short-lived) consumer in period t and market i is given by

$$(q_i^t - 1)b_i^t = (e_{G,i}^t e_{L,i}^t - 1)b_i^t.$$

We seek the perfect public equilibrium (PPE) that maximizes the joint value of G and L_i,

$$V_i \equiv \sum_{t=0}^{\infty} \delta^t \left[b_i^t - (e_{G,i}^t \lambda_G + e_{L,i}^t \lambda_L)c \right],$$

and is symmetric in the sense that all n local players use identical strategies. Let $\overline{V}_i$ and $\underline{V}_i$ denote the maximal and minimal values of V_i, respectively, that can be sustained in a PPE.

Public histories. We model the distinction between independent and collective branding purely as differences in public information concerning the signals s. With independent brands, the public signal in market i is the market-specific signal s_i together with the randomization device r. Consequently, the public history, h_i^t, in market i at time t is

$$h_i^t = (s_i^\tau, r^\tau)_{\tau=0,\ldots,t}.$$

In contrast, the public signal under collective branding consists only of the aggregate signal $\tilde{s}^\tau = \sum_i s_i^\tau$ together with the randomization device r.[14] Consequently, the public history, h^t, in market i at time t is

$$h^t = (\tilde{s}^\tau, r^\tau)_{\tau=0,\ldots,t},$$

where $\tilde{s}^0 \equiv n$, and for $\tau > 0$, $\tilde{s}^\tau$ equals the number of positive realizations of the n noisy quality signals, $\tilde{s}^\tau = \sum_i s_i^\tau$.

We view this difference in public histories as capturing the basic idea in marketing that consumers identify the quality of a good through its brand name alone. Indeed, the idea implies that, under collective branding, consumers cannot discern information about the good's quality that is market specific.[15] The public history can therefore only contain aggregate signals of identically branded products. Our focus on perfect public equilibrium (PPE) then implies that we study only behavior in which players condition their strategies on the coarse brand-specific public history rather than any finer information. This also means that, under

[14] Rather than its sum, we may take the aggregated signal $\tilde{s}^\tau$ as any symmetric and strictly increasing function of the individual signals $(s_1^\tau, \ldots, s_n^\tau)$.

[15] For example, consider a beer tasting website such as *beeradvocate.com*. Even though large beer brands are often produced in several plants, including under license in foreign countries, the tasting notes on such websites do not distinguish between them.

independent branding, players' strategies in market i are independent of the quality signals in another market $j \neq i$.

To ensure that our analysis is non-trivial, we assume throughout that $c < \bar{c} \equiv (1 - \alpha - \beta)/(1 - \beta)$. This assumption is necessary and sufficient for effort to be sustainable for a large enough discount factor under independent branding.

3 Independent Branding

In this section, we analyze equilibrium outcomes when goods in the different markets are branded independently. Since all markets are symmetric and independent, we fix some market i and drop the market subscript for the remainder of this section. All our payoff results are therefore in terms of "per market averages". We sometimes use the superscript I for denoting optimal solutions in this case of independent branding.

Worst PPE. Neither G nor L exerting any effort ($e^t_G = e^t_L = 0$) and short-lived consumers not purchasing the good ($b^t = 0$) in every period t after any history h^t is a PPE. In this PPE, both G and L receive their minmax payoff of zero. This minmax equilibrium outcome represents the worst PPE outcome with the associated payoff of $\underline{V}^I = 0$.

Strategy profiles for best PPE. If the best PPE yields a strictly positive payoff, $\overline{V}^I > \underline{V}^I = 0$, then it involves players exerting effort in equilibrium. Since effort is costly, such an equilibrium must provide players with incentives to induce it. By Abreu, Pearce, and Stacchetti (1990)'s bang-bang result for repeated games with imperfect public monitoring, it is without loss to assume that a PPE takes on only extremal points of the equilibrium value set. Because the most extreme punishments are the minmax payoffs of zero, the extremal points that provide incentives for effort involve a probabilistic triggering of these minmax payoffs that are induced by a market breakdown.

With independent branding, where there is only a binary public signal s on which players can condition their behavior, the description of extremal points that provide incentives by reverting to the minmax outcome is straightforward. They only involve a probability of market breakdown, ρ_0, in case that the signal s points to shirking, i.e., $s = 0$. More specifically, a strategy profile $\sigma^I(\rho_0)$ sustaining effort through the use of extremal points has the following structure: for $\rho_0 \in (0, 1]$, if the period-t history h^t involves $s^\tau = 0$ and $r^\tau \in [0, \rho_0]$ for some $\tau \leq t$, then $e^t_G = e^t_L = 0$ and $b^t = 0$; otherwise, $e^t_G = e^t_L = 1$ and $b^t = 1$.

The strategy profile $\sigma^I(\rho_0)$ implies that, in period 0, both G and L exert effort, and the consumer purchases the good. This continues in all subsequent periods until the public quality signal assumes the value of zero (falsely indicating that the quality in the previous period was zero) *and* the realized value of the public randomization device is not larger than ρ_0; from then on, no effort will ever be exerted and the good is not purchased. That is, a bad quality signal triggers a reversion to the worst PPE with probability ρ_0.[16]

Payoffs and market breakdown probabilities. Playing the strategy profile $\sigma^I(\rho_0)$ yields a long-lived player $j \in \{H, R\}$ a payoff

$$\tilde{V}_j = \pi_j - \lambda_j c + \delta(1-p_0)\tilde{V}_j = \lambda_j V(p_0, \gamma_j), \text{ with } V(p_0, \gamma_j) \equiv \frac{\gamma_j - c}{1 - \delta(1-p_0)},$$

where $\gamma_j \equiv \pi_j / \lambda_j$ is the *reward-to-cost-share ratio*, and $p_0 \equiv \alpha\rho_0$ the expected probability that, in any period after which effort was exerted and the consumer purchased the good, the long-run players stop exerting effort and consumers stop purchasing the good. We refer to p_0 as the on-path *market breakdown probability* when none of the players shirk. $V(p_0, \gamma_j) \geq 0$ if and only if $\gamma_j \geq c$. In this case, $V(p_0, \gamma_j)$ is decreasing in p_0 and increasing in γ_j. Using the identity $\gamma_G = \pi_G/\lambda_G = (1-\lambda_L\gamma_L)/(1-\lambda_L)$, it follows that the payoffs of both long-lived players exceed the minmax payoff of zero if and only if $\gamma_G \in [c, (1-\lambda_L c)/(1-\lambda_L)]$ and $\gamma_L \in [c, (1-\lambda_G c)/(1-\lambda_G)]$.

Incentive constraints. In equilibrium, every consumer receives a payoff of zero, and it is straightforward to see that no consumer has an incentive to deviate from the above strategy profile (which gives him just his minmax payoff of zero). To see whether G or L are better off deviating, note first that the answer is trivially no once the reversion to the worst PPE has been triggered. Consider now a one-shot deviation before such a reversion has been triggered: j's value from one-time shirking is equal to

$$\begin{aligned} \tilde{V}_j^d &= \pi_j + \delta(1-p_1)\tilde{V}_j \\ &= \lambda_j \times [\gamma_j + \delta(1-p_1)V(p_0, \gamma_j)], \end{aligned}$$

where

$$p_1 \equiv (1-\beta)\rho_0$$

is the market breakdown probability in the period after shirking by one of the players.

The incentive constraint for $j \in \{G, L\}$, $\tilde{V}_j \geq \tilde{V}_j^d$, can be written as

$$\delta(p_1 - p_0)V(p_0, \gamma_j) \geq c. \qquad (IC_j^I)$$

[16] Instead of a probabilistic permanent transition to the worst PPE, an alternative strategy profile would involve a deterministic transition to a finite punishment phase of length T, thus not requiring the existence of a public randomization device. In the absence of integer constraints on T, such deterministic strategies would support the same equilibrium outcome.

The left-hand side represents the long-term loss from the one-shot deviation, induced by an increase in the market breakdown probability from p_0 to p_1, whereas the right-hand side represents the short-run gain, which equals the saved effort cost.

Characterizing the best PPE. In the best PPE, the punishment probability ρ_0 maximizes aggregate surplus

$$\tilde{V}_G + \tilde{V}_L = \lambda_G V(\alpha\rho_0, \gamma_G) + \lambda_L V(\alpha\rho_0, \gamma_L) = \frac{1-c}{1-\delta[1-\alpha\rho_0]}$$

subject to (IC_G^I) and (IC_L^I). As $V(\alpha\rho_0, \gamma_j)$ is decreasing in ρ_0, this amounts to minimizing ρ_0 subject to (IC_G^I) and (IC_L^I). Moreover, as (IC_j^I) is violated for ρ_0 small (provided $\lambda_j > 0$), the conditional punishment probability in the best PPE must be such that one of the two incentive constraints holds with equality, and the other with a weak inequality. That is,

$$\rho_0 = \frac{(1-\delta)c}{\delta\left[(1-\alpha-\beta)\min_j \gamma_j - (1-\beta)c\right]}, \tag{1}$$

provided the right-hand side is positive and not larger than one, which holds if and only if[17]

$$\min_j \gamma_j \geq \frac{(1-\beta\delta)c}{\delta(1-\alpha-\beta)}. \tag{2}$$

If (2) does not hold, then an equilibrium with $e_G^0 = e_L^0 = 1$ and $b^0 = 1$ does not exist and, given (γ_H, γ_L), any equilibrium gives both G and L their minmax payoff of zero.

The joint value of G and L is maximized by $\gamma_G = \gamma_L = 1$, i.e., by ex ante agreeing that each long-lived player's reward share is proportional to his cost share: $(\pi_G, \pi_L) = (\lambda_G, \lambda_L)$. Given this optimal revenue sharing scheme, it follows from (1) that the conditional punishment probability equals

$$\rho_0^I \equiv \frac{(1-\delta)c}{\delta[1-\alpha-\beta-(1-\beta)c]}. \tag{3}$$

This probability is positive and not larger than one (meaning that $\sigma^I(\rho_0^I)$ does indeed form an equilibrium) if and only if

$$\delta \geq \frac{c}{1-\alpha-\beta+\beta c} \equiv \bar{\delta}^I. \tag{4}$$

Our parameter restriction at the end of Section 2, $c < \bar{c}$, implies $\bar{\delta}^I < 1$.

We summarize these results in the following proposition:

[17] It is straightforward to verify that if the r.h.s. of (1) is not larger than one it is also positive.

Proposition 1. *If $\delta \geq \bar{\delta}^I$, the best PPE exhibits both long-lived players, G and L, exerting effort, and the consumer purchasing the good, until the beginning of the punishment phase. In this equilibrium, each long-lived player $j \in \{G, L\}$ has a revenue share equal to its cost share (i.e., $\gamma_G = \gamma_L = 1$); the on-path punishment probability is $p_0^I = \alpha \rho_0^I$; and the joint value is equal to $\overline{V}^I = V(p_0^I, 1) > 0 = \underline{V}^I$.*

Otherwise, if $\delta < \bar{\delta}^I$, any PPE involves both G and L shirking in every period, so that $\overline{V}^I = \underline{V}^I = 0$.

In the best PPE, neither the on-path breakdown probability p_0^I nor the joint value $\overline{V}^I$ depend on the effort cost structure (λ_G, λ_L): By optimally sharing the revenue in proportion to the effort cost, the cost structure becomes irrelevant.

The *average per-period payoff* $\overline{v}^I \equiv (1-\delta)\overline{V}^I$ in the best PPE is therefore given by

$$\overline{v}^I = \begin{cases} 0 & \text{if } 0 < \delta < \bar{\delta}^I, \\ 1 - c - \left(\frac{\alpha}{1-\alpha-\beta}\right) c & \text{if } \bar{\delta}^I \leq \delta < 1. \end{cases} \tag{5}$$

For $\delta \geq \bar{\delta}^I$, the average payoff $\overline{v}^I$ is independent of the discount factor, and strictly less than the efficient payoff of $(1-c)$.[18] In the limit as the probability of a "false negative" (α) becomes small, this inefficiency vanishes: $\lim_{\alpha \to 0} \overline{v}^I = 1 - c$. While this inefficiency also decreases as the probability of a "false positive" decreases, it does not vanish in the limit as β becomes small: $\lim_{\beta \to 0} \overline{v}^I = 1 - c/(1-\alpha) < 1 - c$. Finally, note that the critical discount factor $\bar{\delta}^I$ is positively related to both α and β, with $\lim_{(\alpha+\beta) \to 0} \bar{\delta}^I = c$.

Indeed, under perfect monitoring ($\alpha = \beta = 0$), high quality provision in every period is sustainable for $\delta \geq c$ and yields a per-period equilibrium value of $1 - c$. Hence, imperfect monitoring exacerbates the implementation of high quality in two ways. First, for a discount factor $\delta \in [c, \bar{\delta}^I)$, high quality is not sustainable with imperfect monitoring whereas it would be under perfect monitoring. Second, for $\delta \geq \bar{\delta}^I$, high quality in the initial period is sustainable both with perfect and imperfect monitoring, but the equilibrium value is lower with imperfect monitoring, $\bar{v}^I < 1 - c$, as high quality cannot be sustained forever.

It is also instructive to compare the equilibrium outcome with that under "vertical integration" where the same agent chooses e_G^t and e_L^t and, in return, gets all of the revenue from selling the good. As (5) does not depend on λ_L, it is immediate that the values of the best and worst PPEs under vertical integration coincide with those under vertical separation. This confirms that, in our model, the physical team production problem can be solved costlessly.

[18] As is well-known, the folk theorem does not generally apply to symmetric PPE with short-run players.

4 Collective Branding

In this section, we analyze equilibrium outcomes when the long-lived players sell the goods in the n markets under one collective brand. In this case, the public history h^t contains the aggregated signals $\tilde{s}^\tau = \sum_i s_i^\tau$ of the previous periods $\tau \leq t$ rather than the individual signals s_i^τ. Since the public signal $\tilde{s}^\tau$ has $n+1$ possible realizations rather than only two as in the case of independent branding, the players' strategies in a PPE with collective branding are potentially more complex.

Worst PPE. Nevertheless, similar to the case of independent branding, there always exists a PPE in which, irrespective of the public histories, all long-lived players exert no effort in every period, and in each market consumers do not purchase the good. As this gives each long-lived player its minmax payoff of zero, the joint value in the worst PPE under collective branding is $\underline{V}^C = 0$, coinciding with the worst PPE outcome under independent branding, $\underline{V}^I = 0$.

Poisson's binomial distribution. In order to characterize the best equilibrium outcome, note that under collective branding, the aggregated public signal $\tilde{s}$ describes the number of markets in which the quality signal s_i indicated that quality was high. That is, $\tilde{s}$ has $n+1$ possible realizations. The probability distribution of these aggregated signals $\tilde{s}$ depends on the distribution of the underlying market-specific signals s_i.

If players exert effort in all n markets, the aggregated signal $\tilde{s}$ follows the standard binomial distribution of n independent Bernoulli trials, each with the identical success probability $1-\alpha$. However, if shirking occurs in some (but not all) markets, the distribution of $\tilde{s}$ does not correspond to a standard binomial distribution, since the success probability in a market without shirking is $1-\alpha$, whereas it is only β in a market where shirking occurs. In particular, if shirking takes place in k of the n markets, $\tilde{s}$ obtains from n trials of which $n-k$ trials have a success probability of $1-\alpha$, and k trials have a success probability of β.

Consequently, the distribution of $\tilde{s}$ is the convolution of the binomial distribution of k trials with success probability β and the binomial distribution of $n-k$ trials with success probability $1-\alpha$.[19] In particular, the distribution $\mathbb{P}_n(\cdot|k)$ is the convolution of $\mathbb{P}_{n-1}(\cdot|k)$ and a signal from a market without shirking, as well as the convolution of $\mathbb{P}_{n-1}(\cdot|k-1)$ and a signal from a market with shirking. Hence, $\mathbb{P}_n(s|k)$ exhibits the following recursive structure:

$$\mathbb{P}_n(s|k) = (1-\alpha)\mathbb{P}_{n-1}(s-1|k) + \alpha\mathbb{P}_{n-1}(s|k) = \beta\mathbb{P}_{n-1}(s-1|k-1) + (1-\beta)\mathbb{P}_{n-1}(s|k-1). \quad (6)$$

[19]Our analysis does not require the use of an explicit formula for $\mathbb{P}_n(s|k)$. However, for completeness, we report here that, following Rukhin et al. (2009), $\mathbb{P}_n(s|k)$ can be written as $\mathbb{P}_n(s|k) = \sum_{i=0}^{k} \binom{k}{i}\binom{n-k}{s-i}\beta^i(1-\beta)^{k-i}(1-\alpha)^{s-i}\alpha^{n-k-s+i}$, using the convention that the binomial coefficient $\binom{k}{i}$ is 0 for a negative integer i.

Being a special case of "Poisson's binomial distribution" (Wang, 1993), the probability distribution $\mathbb{P}_n(\cdot|k)$ is unimodal and log concave with expectation $\mathbb{E}_n(s|k) = k(1-\beta) + (n-k)\alpha$. A property, crucial for our analysis, is that Poisson's extension of the binomial distribution retains the monotone likelihood ratio property (MLRP). That is, for all $s, k \in \{1, \ldots, n\}$, the following holds:

$$\frac{\mathbb{P}_n(s|k)}{\mathbb{P}_n(s|k-1)} < \frac{\mathbb{P}_n(s-1|k)}{\mathbb{P}_n(s-1|k-1)}.$$

Since MLRP implies FOSD, this also means that the distributions $\mathbb{P}_n(\cdot|k)$ are ordered in the sense of first-order stochastic dominance (FOSD). The FOSD-relation reflects the simple intuition that when shirking in one more market, it is less likely to see at least the same number of successes as without this additional shirking. Yet, even though it is less likely to see at least the same number of successes, the recursive structure (6) implies that it is more likely to see at least one less success with the additional shirking.[20] Hence, for our Poisson's binomial distribution the magnitude of FOSD is also limited: for all $k = 0, \ldots, n-1$ and $s = 1, \ldots, n$, it holds that

$$\mathbb{P}_n(\tilde{s} \geq s|k+1) \leq \mathbb{P}_n(\tilde{s} \geq s|k) \leq \mathbb{P}_n(\tilde{s} \geq s-1|k+1). \tag{7}$$

The first inequality is FOSD, the second inequality describes the sense in which FOSD is limited.

Collective-branding strategies. Similar to the analysis with independent branding, we consider collective-branding strategy profiles $\sigma^C(\cdot)$ that are characterized by $n+1$ punishment probabilities, $\{\rho_s\}_{s=0}^n$, where s indicates the number of positive quality signals.[21] In particular, if the period-t history is such that either $\tilde{s}^\tau = s$ and $r^\tau \in [0, \rho_s]$ for some $\tau \leq t$, then $e^t_{G,1} = e^t_{G,2} = e^t_{L,1} = e^t_{L,2} = 0$ and $b^t_1 = b^t_2 = 0$; otherwise, $e^t_{G,1} = e^t_{G,2} = e^t_{L,1} = e^t_{L,2} = 1$ and $b^t_1 = b^t_2 = 1$.

The strategy profile $\sigma^C(\cdot)$ implies that the repeated game starts with all long-lived players exerting effort, and in each market the consumer purchasing the good. This continues in all subsequent periods until the number of realized positive quality signals in some future period t is $\tilde{s}^t$ and the realization of the public randomization device is less than $\rho_{\tilde{s}^t}$, which then triggers a reversion to the worst PPE.

[20] To see this, note that by (6) it holds, $\mathbb{P}_n(\tilde{s} \geq s-1|k+1) = \beta\mathbb{P}_{n-1}(s-2|k) + \sum_{j=s-1}^{n-1} \mathbb{P}_{n-1}(j|k) + (1-\beta)\mathbb{P}_{n-1}(n|k)$ and $\mathbb{P}_n(\tilde{s} \geq s|k) = (1-\alpha)\mathbb{P}_{n-1}(s-1|k) + \sum_{j=s}^{n-1} \mathbb{P}_{n-1}(j|k) + \alpha\mathbb{P}_{n-1}(n|k)$, where $\mathbb{P}_{n-1}(n|k) = 0$. Subtracting the second from the first yields $\mathbb{P}_n(\tilde{s} \geq s-1|k+1) - \mathbb{P}_n(\tilde{s} \geq s|k) = \beta\mathbb{P}_{n-1}(s-2|k) + \alpha\mathbb{P}_{n-1}(s-1|k) \geq 0$.

[21] Since $\tilde{s} = n$ is indicative of no-shirking, it will always be optimal to have $\rho_n = 0$.

Market breakdown probabilities. Conditional on all long-lived players having exerted effort in the past, and consumers having purchased the goods, the strategy profile $\sigma^C(\cdot)$ induces breakdown probabilities, both on-path as well as following a deviation. The breakdown probability in the period after shirking in k markets is denoted p_k, and given by

$$p_k = \sum_{s=0}^{n} \mathbb{P}_n(s|k)\rho_s. \tag{8}$$

Incentive constraints. Averaging across the n markets, G's average value, $\tilde{V}_G$, when players play the strategy profile $\sigma^C(\cdot)$ equals $\pi_G - \lambda_G c + \delta(1-p_0)\tilde{V}_G$, implying

$$\tilde{V}_G = \frac{\pi_G - \lambda_G c}{1-\delta(1-p_0)} = \lambda_G V(p_0, \gamma_G). \tag{9}$$

The global player is free to choose different effort levels in different markets.[22] Her (average-per-market) value from shirking in k markets in the current period and subsequently reverting to the collective branding strategy $\sigma^C(\cdot)$ is

$$\tilde{V}_G^{d,k} = \pi_G - \frac{n-k}{n}\lambda_G c + \delta[1-p_k]\tilde{V}_G, \tag{10}$$

where the second term on the right-hand side represents the average-per-market effort cost when shirking in k markets and exerting effort in the other $n-k$ markets. We can rewrite the incentive constraint, $\tilde{V}_G \geq \tilde{V}_G^{d,k}$, as

$$\delta(p_k - p_0)V(p_0, \gamma_G) \geq c \cdot \frac{k}{n}. \tag{$IC_G^{C,k}$}$$

Intuitively, the left-hand side represents the (average) long-term loss — the breakdown probability rising from p_0 to p_k — from the one-shot deviation, whereas the right-hand side represents the average short-run gain — the (per-market-average) reduction in effort costs — from that deviation.

Dropping the market subscript for notational convenience, local player L's value, $\tilde{V}_L$, when players play the strategy profile $\sigma^C(\cdot)$ equals $\pi_L - \lambda_L c + \delta(1-p_0)\tilde{V}_L$, implying

$$\tilde{V}_L = \frac{\pi_L - \lambda_L c}{1-\delta(1-p_0)} = \lambda_L V(p_0, \gamma_L). \tag{11}$$

The local player's value from the one-shot deviation to shirking in the current period and

[22] In the conclusion, we briefly discuss the case where G has to take the same action in all markets.

subsequently reverting back to the collective branding strategy $\sigma^C(\cdot)$ is

$$\tilde{V}_L^d = \pi_L + \delta(1 - p_1)\tilde{V}_L. \tag{12}$$

We can write the local player's incentive-constraint, $\tilde{V}_L \geq \tilde{V}_L^d$, as

$$\delta(p_1 - p_0)V(p_0, \gamma_L) \geq c. \tag{IC_L^C}$$

Intuitively, the left-hand side represents the long-term loss — the breakdown probability rising from p_0 to p_1 — from the one-shot deviation, whereas the right-hand side represents the short-run gain — the reduction in effort costs — from that deviation.

Characterizing the best PPE. The best PPE is characterized by the vector of punishment probabilities $\rho^C = (\rho_0^C, \ldots, \rho_n^C)$ that maximizes aggregate surplus

$$n\tilde{V}_G^d + n\tilde{V}_L^d = \frac{n(1-c)}{1 - \delta(1-p_0)} = nV(p_0, 1) \tag{13}$$

subject to the $(n+1)$ incentive constraints $(IC_G^{C,k})$ and (IC_L^C). Note that maximizing aggregate surplus is equivalent to minimizing the on-path breakdown probability p_0, which is linear in the punishment probabilities ρ_s. Because we can also express the incentive constraint as constraints that are linear in the punishment probabilities ρ_s, characterizing the best PPE under collective branding involves solving a linear programming problem with $n+1$ constraints. In particular, ρ^C is a solution to the linear program

$$\begin{aligned}
\mathcal{P}^n : \min_{(\rho_0,\ldots,\rho_n)} \quad & \sum_{s=0}^{n} \mathbb{P}_n(s|0)\rho_s \\
\text{s.t.} \quad & \sum_{s=0}^{n} \left[\frac{\mathbb{P}_n(s|k) - \mathbb{P}_n(s|0)}{k} - \frac{c\mathbb{P}_n(s|0)}{n(\gamma_G - c)}\right]\rho_s \geq \frac{(1-\delta)c}{n\delta(\gamma_G - c)}, \;\forall\, k \in \{1, \ldots, n\} \\
& \sum_{s=0}^{n} \left[\mathbb{P}_n(s|1) - \mathbb{P}_n(s|0) - \frac{c\mathbb{P}_n(s|0)}{\gamma_L - c}\right]\rho_s \geq \frac{(1-\delta)c}{\delta(\gamma_L - c)}.
\end{aligned}$$

The next lemma is our main step towards a characterization of solutions to $\mathcal{P}$.

Lemma 1. *Suppose ρ^C is a solution to $\mathcal{P}$. Then for ρ^C at least one constraint is binding and there is an integer $\bar{s} < n$ such that $\rho_s^C = 1$ for all $s < \bar{s}$ and $\rho_s^C = 0$ for all $s > \bar{s}$. The solution implies $p_0 \leq \ldots \leq p_n$.*

The lemma shows that, optimally, the punishment probabilities are concentrated in a bang-bang fashion on the lowest values of the public aggregated signal $\tilde{s}$. In particular, there is a *cutoff signal* $\bar{s}$ such that all realization of $\tilde{s}$ that lie below this threshold imply a market breakdown with certainty, whereas realizations of $\tilde{s}$ above the threshold imply no market breakdown whatsoever. Formally, the result follows from the MLRP property of Poisson's binomial distribution, which reflects the intuitive notion that low values of the public signal $\tilde{s}$ are less likely when agents put in effort. Intuitively, this property implies that concentrating the punishment probabilities on the lowest values of $\tilde{s}$ provides the strongest incentives for effort.

An algorithm. Lemma 1 allows the construction of an algorithm for finding a solution to $\mathcal{P}^n$. First, define $\bar{s}$ as the smallest $\tilde{s} \in \{0, \ldots, n-1\}$ such that

$$\sum_{s=0}^{\tilde{s}} \Delta_G(k, s) \geq \frac{(1-\delta)c}{n\delta(\gamma_G - c)}, \qquad \forall k = 1, ..., n, \tag{14}$$

and

$$\sum_{s=0}^{\tilde{s}} \Delta_L(1, s) \geq \frac{(1-\delta)c}{\delta(\gamma_L - c)}, \tag{15}$$

where

$$\Delta_G(k, s) \equiv \frac{\mathbb{P}_n(s|k) - \mathbb{P}_n(s|0)}{k} - \frac{c\mathbb{P}_n(s|0)}{n(\gamma_G - c)} \text{ and } \Delta_L(k, s) \equiv \frac{\mathbb{P}_n(s|k) - \mathbb{P}_n(s|0)}{k} - \frac{c\mathbb{P}_n(s|0)}{\gamma_L - c}.$$

The variable $\bar{s}$ is found algorithmically by starting with $\tilde{s} = 0$ and increasing it successively until either all of the $n+1$ inequalities associated with (14) and (15) hold, or $\tilde{s} = n$. If this procedure ends with $\tilde{s} = n$, then $\bar{s}$ does not exist, implying that the feasible set of $\mathcal{P}^n$ is empty so that there is no equilibrium in which agents supply effort. If the procedure ends with $\tilde{s} < n$, then $\bar{s} = \tilde{s}$. In a next step, compute for $k = 1, \ldots, n$,

$$\rho_G^k \equiv \frac{1}{\Delta_G(k, \bar{s})} \left\{ \frac{(1-\delta)c}{n\delta(\gamma_G - c)} - \sum_{s=0}^{\bar{s}-1} \Delta_G(k, s) \right\}; \quad \rho_L \equiv \frac{1}{\Delta_L(1, \bar{s})} \left\{ \frac{(1-\delta)c}{\delta(\gamma_L - c)} - \sum_{s=0}^{\bar{s}-1} \Delta_L(1, s) \right\}.$$

By construction, each ρ_G^k and ρ_L lies in $[0, 1]$. Taking $\bar{\rho}$ as the maximum over all ρ_G^k and ρ_L, it then follows from Lemma 1 that the solution ρ^C of program $\mathcal{P}^n$ exhibits $\rho_s^C = 1$ for $s < \bar{s}$, $\rho_{\bar{s}}^C = \bar{\rho}$, and $\rho_s^C = 0$ for $s > \bar{s}$, with an attained objective of $p_0 = \sum_{s=0}^{\bar{s}} \mathbb{P}_n(s|0) + \mathbb{P}_n(\bar{s}|0)\bar{\rho}$.

The algorithm is useful for identifying the comparative statics of the threshold signal $\bar{s}$ with respect to the discount factor δ. To see this, first note that the left-hand sides of the $n+1$ inequalities associated with (14) and (15) are independent of δ, whereas the right-hand sides decrease with δ over $[0, 1]$ and grow arbitrarily large as δ approaches 0. This implies that the threshold $\bar{s}$ fails to exist when δ is small and when it does exist for some $\hat{\delta}$, it exists for all $\delta > \hat{\delta}$ and, moreover, $\bar{s}$ is decreasing in δ.

Defining

$$\bar{\gamma}_G \equiv c + \frac{c\alpha}{n(1-\alpha-\beta)} \leq \bar{\gamma}_L \equiv c + \frac{c\alpha}{1-\alpha-\beta},$$

we collect and extend these insights in the following lemma.

Lemma 2. *Suppose $\gamma_G > \bar{\gamma}_G$ and $\gamma_L > \bar{\gamma}_L$. Then there is a critical discount factor $\bar{\delta}^C \in [0,1)$ such that the threshold value $\bar{s}$ exists if and only if $\delta \geq \bar{\delta}^C$. For $\delta > \bar{\delta}^C$, the cut-off signal $\bar{s}$ is decreasing in δ. In particular, there is a cutoff $\bar{\delta}_0^C < 1$ such that for $\delta > \bar{\delta}_0^C$, we have $\bar{s} = 0$ and, moreover, of the first n constraints in program $\mathcal{P}^n$ at most the constraint w.r.t. to $k = 1$ is binding.*

Collective vs. Independent Branding. In order to identify the effect of a collective reputation, it is helpful to consider first two polar cases: The one in which only the global player, G, has to incur costly effort (i.e., $\lambda_G = 1$), and the one in which only local players have to incur costly effort (i.e., $\lambda_L = 1$). Below, we show that, in the first polar case, collective branding permits sustaining a better reputation and a higher value in the best PPE than independent branding. In the second case, by contrast, collective branding induces free-riding on the other local players' reputation and tends to reduce the maximum sustainable value in the best PPE. In the generic case in which both the global and local players have to exert costly effort, collective branding is beneficial only if each local player receives a revenue share exceeding his cost share so as to mitigate the free-rider effect.

Global Effort Cost Only ($\lambda_L = 0$). In order to show that without any reputational free-riding, collective branding is optimal, we first study the polar case in which the global player incurs all effort costs, i.e., $\lambda_G = 1$. Since $\lambda_G = 1$ implies $\lambda_L = 1 - \lambda_G = 0$, the local players' effort in this polar case is costless so that they do not need any incentives to exert effort. It is therefore optimal to give the entire revenue share of each good to G, implying that, just as under independent branding, the proportional reward scheme, $(\pi_G, \pi_L) = (\lambda_G, \lambda_L)$, is optimal.

First consider the case $\rho_0^I \leq 1$ so that under independent branding, the best PPE exhibits $\overline{V}^I > 0$ and minimizes the on-path market breakdown probability, p_0, with respect to the punishment probability, ρ_0, subject to the incentive constraint (IC_G^I).

Given $\overline{V}^I > 0$, a first natural question to ask is whether collective branding can also attain this value. If so, then we must have $\overline{V}^C \geq \overline{V}^I > 0$ so that the best PPE under collective branding minimizes p_0 with respect to $\rho = (\rho_1, \ldots, \rho_n)$ subject to the n incentive constraints ($IC_G^{C,k}$).

We next argue that collective branding can indeed attain the value $\overline{V}^I > 0$ and replicate the outcome under independent branding. To see this, consider the following punishment probabilities under collective branding:

$$\rho_s = \rho_s(\rho_0^I) \equiv \frac{n-s}{n}\rho_0^I, \qquad s = 0, ..., n.$$

Since $\rho_s(\rho_0^I) \leq \rho_0^I$, it follows from $\overline{V}^I > 0$ that $\rho_0^I \in [0,1]$ and therefore $\rho_s(\rho_0^I) \in [0,1]$ for all $s = 0, ..., n$. Moreover, the punishment probabilities $\rho_s = \rho_s(\rho_0^I)$ lead to a market breakdown after shirking in k markets of

$$p_k = \sum_{s=0}^{n} \mathbb{P}_n(s|k)\rho_s = \sum_{s=0}^{n} \mathbb{P}_n(s|k)\frac{n-s}{n}\rho_0^I = \frac{(n - \mathbb{E}[s|k])}{n}\rho_0^I = \frac{(n-k)\alpha + k(1-\beta)}{n}\rho_0^I.$$

In particular, the on-path breakdown probability p_0 coincides with the one under independent branding: $p_0 = \alpha\rho_0^I$. Moreover, for these punishment probabilities each $(IC_G^{C,k})$ is equivalent to the incentive constraint (IC_G^I):

$$\frac{n}{k} \cdot \delta(p_k - p_0)V(p_0, 1) \geq c \;\Leftrightarrow\; \delta\rho_0^I(1 - \alpha - \beta)V(p_0, 1) \geq c.$$

Hence, by taking $\rho_s = \rho_s(\rho_0^I)$ we can replicate the best PPE under independent branding. However, since $\rho_s(\rho_0^I)$ violates the optimal cutoff structure (except for the special case of $n = 2$ and $\rho_0^I = 1$), Lemma 1 implies that we can strictly improve upon this outcome. It follows that the best PPE yields a strictly higher payoff under collective branding. This leads to the following result:

Proposition 2. *Suppose $\lambda_G = 1$. Then the optimal rewards exhibit $\hat{\gamma}_G = 1$. Moreover, if $\rho_0^I \leq 1$, then collective branding is superior to independent branding and $\bar{\delta}^C \leq \bar{\delta}^I$. This superiority is strict except for the special case $n = 2$ and $\rho_0^I = 1$. If $\rho_0^I > 1$, then collective branding is superior to independent branding and strictly so for $n > 2$ and ρ_0^I close to one.*

The proposition identifies the beneficial informativeness effect of collective branding. It shows that collective branding increases payoffs by concentrating market breakdown on events with a sufficiently large number of bad quality signals. By pooling these signals and punishing the global player (only) if a sufficiently large number of them is bad, the on-path breakdown probability p_0 can be reduced, thereby mitigating the inefficiency caused by imperfect monitoring. This informativeness effect is an implication of the natural MLR property of Poisson's binomial distribution: Lemma 1's optimality result – showing that only punishment probabilities that display a cutoff structure use the collective signal optimally – is based on that property. We therefore conclude that collective branding is optimal, even though signals are independent across markets and the global player has the flexibility to shirk in any number of markets.

Local Effort Cost Only ($\lambda_L = 1$). We next turn to the other polar case in which only local players incur effort costs, i.e., $\lambda_L = 1$. As this implies $\lambda_G = 1 - \lambda_L = 0$, the global player's effort can be induced "for free" in that G does not need any incentives to exert effort. Hence, it is optimal to give the entire revenue share of each good i to L_i, implying once more that the proportional reward scheme, $(\pi_G, \pi_L) = (\lambda_G, \lambda_L)$, is optimal.

Similar to analyzing the previous polar case, we first consider the case $\rho_0^I \leq 1$ so that under independent branding, the best PPE exhibits $\overline{V}^I > 0$ and ρ_0^I minimizes the on-path market breakdown probability, $p_0 = \alpha\rho_0^I$, subject to the incentive constraint (IC_L^I).

Given $\overline{V}^I > 0$, we again ask the question whether collective branding can also attain this value. If so, then we must have $\overline{V}^C \geq \overline{V}^I > 0$ so that the best PPE under collective branding minimizes p_0 with respect to $\rho = (\rho_0, \dots, \rho_n)$ subject to the incentive constraint (IC_L^C). More precisely, $\rho^C = (\rho_0^C, \dots, \rho_n^C)$ now solves

$$\begin{aligned}
\mathcal{P}^L : \min_{(\rho_0,\dots,\rho_n)} \quad & \sum_{s=0}^{n} \mathbb{P}_n(s|0)\rho_s \\
\text{s.t.} \quad & \sum_{s=0}^{n} \left[\mathbb{P}_n(s|1) - \mathbb{P}_n(s|0) - \frac{c\mathbb{P}_n(s|0)}{1-c}\right]\rho_s \geq \frac{(1-\delta)c}{\delta(1-c)}.
\end{aligned}$$

Problem $\mathcal{P}^L$ is a relaxed version of $\mathcal{P}^n$ since it lacks the first n constraints. Noting that Lemma 1 applies also with respect to the more relaxed problem $\mathcal{P}^L$, it follows that ρ^C displays a cutoff structure with some cut-off signal $\bar{s}$. The minimum cutoff is $\bar{s} = 0$. This is indeed the optimal cutoff if and only if

$$\left[\mathbb{P}_n(0|1) - \mathbb{P}_n(0|0) - \frac{c\mathbb{P}_n(0|0)}{1-c}\right] \geq \frac{(1-\delta)c}{\delta(1-c)}. \tag{16}$$

In this case, $\rho_s^C = 0$ for $s > 0$ and, using $\mathbb{P}_n(0|1) = \alpha^{n-1}(1-\beta)$ and $\mathbb{P}_n(0|0) = \alpha^n$,

$$\rho_0^C = \bar{\rho}_0^L \equiv \frac{(1-\delta)c}{\delta\alpha^{n-1}[1-\alpha-\beta-(1-\beta)c]} = \frac{\rho_0^I}{\alpha^{n-1}}.$$

Hence, (16) is equivalent to $\rho_0^I \leq \alpha^{n-1}$. It follows that, for $\rho_0^I \leq \alpha^{n-1}$, the aggregate surplus associated with the best PPE under collective branding matches the aggregate surplus associated with the best PPE under independent branding so that we have $\overline{V}^C = \overline{V}^I$.

Note that if $\rho_0^I > \alpha^{n-1}$, then the best PPE with collective branding must be strictly worse than with independent branding, since either collective branding can implement effort only with punishment probabilities $\rho_0^C = 1$ and $\rho_1^C > 0$, which due to MLRP, yields a larger breakdown probability p_0 — implying $\overline{V}^I > \overline{V}^C > 0$, or collective branding cannot implement high effort — implying $\overline{V}^I > \overline{V}^C = 0$. In either case, we have $\overline{V}^I > \overline{V}^C$, implying that collective branding does strictly worse. This also means that $\overline{V}^I = 0$ implies $\overline{V}^C = 0$ so that the critical discount factor at which effort is sustainable with collective branding cannot be smaller than the corresponding discount factor under independent branding, $\bar{\delta}^I \leq \bar{\delta}^C$. We collect these insights in the following proposition.

Proposition 3. *Suppose $\lambda_L = 1$. Then, independent branding is superior to collective branding, i.e., $\overline{V}^I \geq \overline{V}^C$, and $\bar{\delta}^I \leq \bar{\delta}^C$. This superiority is strict whenever $\rho_0^I \in (\alpha^{n-1}, 1]$. For $\rho_0^I \leq \alpha^{n-1}$ the superiority is weak in that $\overline{V}^I = \overline{V}^C$.*

The proposition identifies the harmful free-rider effect of collective branding. To see this effect more clearly, recall that – under both independent and collective branding – the local player's incentive constraint can be written as

$$\delta \left(\frac{p_1}{p_0} - 1 \right) p_0 V(p_0, \gamma_L) \geq c,$$

where $\gamma_L = 1$ in the best PPE. Under independent branding, $p_1/p_0 = (1-\beta)/\alpha > 1$. By contrast, under collective branding, we can exploit the recursive structure (6) to rewrite p_0 and p_1 associated with $\rho^C = (1, \ldots, 1, \rho_{\bar{s}}^C, 0, \ldots, 0)$ with a cutoff $\bar{s}$ as follows:

$$p_0 = \alpha\Delta + B; \text{ and } p_1 = (1-\beta)\Delta + B,$$

where

$$\Delta \equiv \mathbb{P}_{n-1}(\bar{s} - 1|0)(1 - \rho_{\bar{s}}^C) + \mathbb{P}_{n-1}(\bar{s}|0)\rho_{\bar{s}}^C > 0;$$

and

$$B \equiv \mathbb{P}_n(\tilde{s} \leq \bar{s} - 1|0) + \mathbb{P}_{n-1}(\bar{s} - 1|0)\rho_{\bar{s}}^C \geq 0.$$

Hence, under collective branding and for any ρ^C exhibiting a cutoff structure, the ratio of punishment probabilities satisfies

$$\frac{p_1}{p_0} = \frac{(1-\beta)\Delta + B}{\alpha\Delta + B} \geq \frac{1-\beta}{\alpha}, \tag{17}$$

where the inequality is strict if and only if $B = 0$, or equivalently $\bar{s} = 0$. That is, under collective branding, a local player choosing to shirk does not increase the punishment probability by as much as he would under independent branding as he correctly anticipates the other local players to put in effort (and thus likely to generate positive signals). Only in the case in which, under collective branding, the transition to the worst PPE occurs only if all signals are bad ($\bar{s} = 0$) is the punishment probability ratio p_1/p_0 the same as under independent branding.

At a more general level, under collective branding, the local player's continuation payoff depends on the signals generated by other players. As he cannot affect those other signals, collective branding can only hurt incentives. Perhaps surprisingly, however, collective branding does as well as independent branding if $\rho_0^I \leq \alpha^{n-1}$. To understand this, note that – under independent branding – the local player is punished only if he generates a bad signal, in which case the transition depends on the realization of the public randomization device. For $\rho_0^I \leq \alpha^{n-1}$, the same on-path punishment probability can be generated under collective branding by transiting to the worst PPE only if all n signals are bad. That is, only if the local player himself as well as all the other $n-1$ local players generate bad signals, implying that $\bar{s} = 0$ so that (17) holds with equality. From the viewpoint of the local player, the outcome of the other signals is purely random and the probability that all of them are bad (given that the other local players do not shirk) equals α^{n-1}. In other words, the randomness of the other $n-1$ signals under collective branding plays the same role as the public randomization device under independent branding and therefore collective branding does not distort incentives. By contrast, if $\rho_0^I > \alpha^{n-1}$, then to generate the same on-path punishment probability under collective branding requires that the transition to the worst PPE may have to occur even if at least one of the n signals is positive. That is, if $\rho_0^I > \alpha^{n-1}$, then $\bar{s} > 0$ so that the inequality in (17) is strict. Since this means that, with some probability, a local player is "punished" even after generating a positive signal, collective branding is strictly harmful for incentives.

Figure 1 contrasts the two polar cases, displaying the comparative statics for the average per period payoff $\overline{v}^C$ and $\overline{v}^I$ with respect to the discount factor δ for those cases. Panel (a) depicts the case where only the global player needs to be incentivized ($\lambda_G = 1$). It illustrates the result of Proposition 2 that collective branding is superior to independent branding in general, and strictly so in two ways. First, in the case $\delta \geq \overline{\delta}^I$, where (initial) effort is sustainable with independent branding, collective branding can sustain it with a strictly lower on-path market breakdown probability (except in the special case $n = 2$ and $\rho_0^I = 1$). Second, in the case $\delta < \overline{\delta}^I$, where (initial) effort is not sustainable with independent branding, collective branding can sustain it for δ smaller but close enough to $\bar{\delta}^I$.

By contrast, Panel (b) illustrates the implied comparative statics for the other polar case, where only the local players have to be incentivized ($\lambda_L = 1$). Defining $\bar{\delta}_L$ as the value of δ at which $\rho_0^I = \alpha^{n-1}$, the values $\overline{V}^C$ and $\overline{V}^I$ coincide for $\delta \geq \bar{\delta}_L$. As illustrated in Panel (b), this implies that, in addition to $\overline{v}^I$, also the maximum average per period payoff under collective branding, $\overline{v}^C$, is constant. For $\delta < \bar{\delta}_L$, we have $\overline{V}^C < \overline{V}^I$ and, due to continuity of $\overline{v}^C$ for $\delta > \bar{\delta}^C$, the maximum average per-period payoff $\overline{v}^C$ is therefore strictly decreasing in the interval $[\bar{\delta}^C, \bar{\delta}_L]$. Moreover, in the interval $[\bar{\delta}^I, \bar{\delta}^C]$ high effort is only sustainable for independent branding. Consequently, the blue curve $\overline{v}^C$ lies always (weakly) below the red curve $\overline{v}^I$ – in stark contrast to Panel (a). In short, Panel (a) displays the

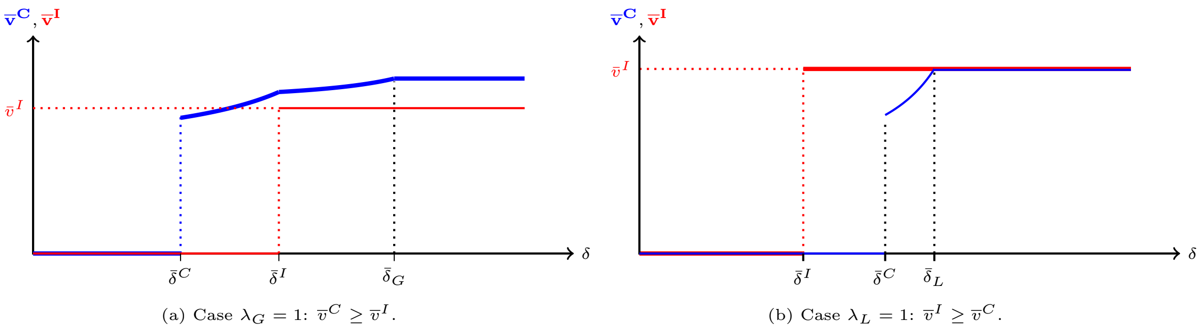

Figure 1: Best PPE's per-period values under independent vs. collective branding.

optimality of collective branding for the case $\lambda_G = 1$, whereas Panel (b) shows the optimality of independent branding for the other polar case $\lambda_L = 1$.

The Generic Case $(\lambda_G \in (0, 1))$**.** We now turn to the generic case in which the global and local players share the overall effort cost c according to the proportions $\lambda_G \in (0, 1)$ and $\lambda_L = 1 - \lambda_G$.

In Section 3, we showed that – under independent branding – it is optimal to provide the long-lived agents with proportional rewards: $(\pi_G, \pi_L) = (\lambda_G, \lambda_L)$. Trivially, this was also the case in the two polar cases of collective branding studied above.

It is therefore instructive to start our analysis of collective branding in the generic case assuming such proportional rewards so that $\gamma_G = \gamma_L = 1$. It then follows that the incentive constraints coincide with the two polar cases studied above. Under collective branding and proportional rewards, the optimal on-path breakdown probability p_0 is minimized subject to

$$\delta(p_k - p_0)V(p_0, 1)\frac{n}{k} \geq c, \qquad k = 1, ..., n; \tag{IC_G^{Ck}}$$

$$\delta(p_1 - p_0)V(p_0, 1) \geq c. \tag{IC_L^C}$$

Note however that by Lemma 1, we have $p_k - p_0 \geq p_1 - p_0$, which together with $k \leq n$ implies that (IC_G^{Ck}) follows from (IC_L^C). As a result, all (IC_G^{Ck}) are redundant so that the optimal on-path breakdown probability p_0 is minimized subject to only (IC_L^C). This however implies that, for proportional rewards $\gamma_G = \gamma_L = 1$, the intermediate case $\lambda_G \in (0, 1)$ boils down to the polar case with only local effort costs $(\lambda_L = 1)$. As a result, Proposition 3 applies, meaning that it extends to all $\lambda_L \in (0, 1]$ and proportional rewards $\gamma_L = 1$:

Proposition 4. *Suppose $\lambda_L \in (0, 1]$. Then collective branding is suboptimal with proportional rewards $(\pi_G, \pi_L) = (\lambda_G, \lambda_L)$, and strictly so for $\rho_0^I \in (\alpha^{n-1}, 1]$.*

This result demonstrates in an extreme sense the drawback of a collective reputation.

As soon as explicit incentives for the local players' effort are needed, independent branding always outperforms collective branding with proportional rewards.

The proposition raises the question whether the long-lived agents can use the reward structure as a tool to mitigate this extreme effect of a collective reputation. We will argue that they can indeed do so; by carefully calibrating the revenue shares γ_G and γ_L, the long-lived players can mitigate the local players' free-riding problem.

To see this, recall that, in general, the incentive constraints depend on the revenue shares γ_G and γ_L as follows:

$$\delta(p_k - p_0)V(p_0, \gamma_G)\frac{n}{k} \geq c, \qquad k = 1, ..., n; \qquad (IC_G^{Ck}(\gamma_G))$$

$$\delta(p_1 - p_0)V(p_0, \gamma_L) \geq c. \qquad (IC_L^C(\gamma_L))$$

Since $V(p_0, \gamma_L)$ is increasing in γ_L, an increase in γ_L relaxes the constraint (IC_L^C). Of course, an increase in the local players' revenue share γ_L is accompanied by a decrease in the global player's revenue share γ_G.[23] Yet, for $\gamma_L = \gamma_G = 1$, each (IC_G^{Ck}) holds strictly whenever the constraint (IC_L^C) binds so that by continuity a small increase in γ_L ensures that the corresponding small decrease in γ_G is such that each (IC_G^{Ck}) is still satisfied, while (IC_L^C) is relaxed. This reasoning suggests that, starting with proportional rewards, we can improve the objective under collective branding by relaxing (IC_L^C) through increasing the local players' reward γ_L. Defining

$$\tilde{\gamma}_G \equiv \frac{1 + (n-1)(1-\lambda_G)c}{\lambda_G + (1-\lambda_G)n} < 1 \text{ and } \tilde{\gamma}_L \equiv \frac{n - (n-1)\lambda_G c}{\lambda_G + (1-\lambda_G)n} > 1,$$

the following lemma refines this intuition and determines the bounds on the optimal γ_G and γ_L.

Lemma 3. *Suppose $\lambda_G \in (0,1)$ and $\overline{V}^C > 0$. Then the optimal reward-to-cost-share ratios $\hat{\gamma}_G$ and $\hat{\gamma}_L$ exhibit $\hat{\gamma}_G \in [\tilde{\gamma}_G, 1)$ and $\hat{\gamma}_L \in (1, \tilde{\gamma}_L]$. In particular, $(\hat{\gamma}_G, \hat{\gamma}_L) = (\tilde{\gamma}_G, \tilde{\gamma}_L)$ for $\delta \in (\bar{\delta}_0^C, 1)$.*

Proposition 4 shows that, with proportional rewards, collective branding is never strictly optimal when incentives are needed for inducing effort from local players. That is, with proportional rewards, the minimum on-path breakdown probability that sustains high quality is lower under independent branding than under collective branding. This however does not automatically imply $\overline{V}^I \geq \overline{V}^C$, since Lemma 3 shows that, when incentives are needed for inducing local players' effort, proportional rewards are never optimal under collective branding. In particular, the lemma leaves open the possibility that, for the optimal reward

[23] In particular, $\gamma_G = \pi_G/\lambda_G = (1-\pi_L)/(1-\lambda_L) = (1-\lambda_L\gamma_L)/(1-\lambda_L)$.

structure, the minimum on-path breakdown probability that sustains high quality is actually lower under collective branding than under independent branding, implying $\overline{V}^C > \overline{V}^I$. The next proposition shows that this is indeed the case for large discount factors—no matter how small the global player's share of the effort cost is, provided it is strictly positive.

Proposition 5. *There exists a cutoff $\overline{\delta} < 1$ such that $\overline{V}^C > \overline{V}^I$ for all $\delta \geq \overline{\delta}$ and all $\lambda_G > 0$.*

The previous results characterize properties of the best PPE without characterizing it explicitly. More specifically, the best PPE has a value $\overline{V}^C > 0$ if and only if the following program has a solution $\hat{p}_0$:

$$\begin{aligned}
\mathcal{P}^C : \min_{(\gamma_G,\gamma_L,\rho_0,\ldots,\rho_n)} \quad & p_0 = \sum_{s=0}^{n} \mathbb{P}_n(s|0)\rho_s \\
\text{s.t.} \quad & \delta(p_k - p_0)V(p_0,\gamma_G)n/k \geq c \quad \forall k = 1,\ldots,n; && (IC_G^{Ck}(\gamma_G)) \\
& \delta(p_1 - p_0)V(p_0,\gamma_L) = c; && (IC_L^C(\gamma_L)) \\
& \gamma_L = (1-\lambda_G\gamma_G)/(1-\lambda_G), && (18)
\end{aligned}$$

where (18) is the accounting identity that links γ_G and γ_L. We denote a solution to $\mathcal{P}^C$ by a triple $(\hat{\gamma}_G, \hat{\gamma}_L, \hat{\rho}) \in \mathbb{R}^{n+1}$. In the case in which $\mathcal{P}^C$ admits a solution $(\hat{\gamma}_G, \hat{\gamma}_L, \hat{\rho})$ with value $\hat{p}_0$, it follows that

$$\overline{V}^C = \frac{1-c}{1-\delta(1-\hat{p}_0)} > 0.$$

Writing out the program that determines the value $\overline{V}^C$ allows us to obtain the following comparative static result:

Proposition 6. *Suppose that $\overline{V}^C > 0$. Then, the optimal reward-to-cost-share ratio $\hat{\gamma}_L$ is strictly increasing, and the induced value $\overline{V}^C$ strictly decreasing, in the cost share λ_L.*

The proposition confirms the intuition that we can interpret λ_L as a measure of the relative magnitude of the free-rider effect. The larger is the share of the effort cost that needs to be borne by the local players, the larger are those players' incentives to take a free ride, and therefore the larger is the reward-to-cost-share ratio $\hat{\gamma}_L$ that the local players optimally receive. While this mitigates the free-rider problem, the flip side is that this decreases the reward-to-cost-share ratio $\hat{\gamma}_G$ and therefore exacerbates the global player's incentive problem so that the resulting aggregate payoff is reduced.

5 Optimal Collective Brand Size

In the previous section, we have shown that, in the presence of local players, a careful calibration of the revenue shares is essential for building a collective brand reputation. In this section, we identify a second essential tool for managing collective branding reputation: the size of the collective brand. In particular, we show that the severity of the moral hazard problem, measured by the parameter λ_L, determines how the sustainability of high quality varies with brand size.

In this section, we thus analyze the comparative statics in the brand size n, obtaining results on the maximum implementable brand size, $\bar{n}$, and the optimal brand size, $\hat{n}$. Our main result reveals that, even when brand size is chosen optimally (without any additional constraints), there remains an inefficiency in the limit as the discount factor becomes large. Despite the usual intractability of Poisson's binomial distributions (e.g., Biscarri et al., 2018), the recursive structure (6) of the probability distribution $\mathbb{P}_n(\cdot|k)$ enables us to express the limiting inefficiency in closed form as a function only of (i) the corresponding limiting inefficiency in the polar case in which the local players bear all effort costs, and (ii) the effort cost share λ_L of the local players.

To make clear the dependencies on the collective brand size n, we denote variables with a superscript n throughout this section. For instance, $\overline{V}^n$ now denotes the average per-market value in the best equilibrium when the brand size is n.

Global Effort Cost Only. ($\lambda_L = 0$) We first consider the benchmark in which collective reputation does not exhibit a free-rider problem. For this benchmark, we show that, for any discount factor $\delta > c$, we obtain efficiency in the limit as n goes to infinity, implying that both the maximum implementable brand size, $\bar{n}$, and the optimal brand size, $\hat{n}$, are unbounded. An intuition behind this efficiency results follows from the observation that with perfect monitoring ($\alpha = \beta = 0$) efficiency obtains if and only if $\delta \geq c$. Extending the brand size mitigates the inefficiency induced by imperfect monitoring. This inefficiency vanishes in the limit as n becomes large.[24]

Recall that this benchmark corresponds to the polar case $\lambda_G = 1$. For this case, the incentive constraint ($IC_L^n(\gamma_L)$) is redundant so that the optimal pair $(\hat{\rho}^n, \hat{\gamma}_G^n)$ that characterizes the best PPE exhibits

$$\hat{\gamma}_G^n = 1 \text{ and } \hat{\rho}^n = \arg\min_{\rho} p_0^n(\rho) \text{ s.t. } (IC_G^{n;1}(1)), \ldots, (IC_G^{n;n}(1)).$$

Proposition 7. *Suppose $\lambda_G = 1$ and $\delta > c$. Then the maximum implementable brand size, $\bar{n}$, and the optimal brand size, $\hat{n}$, are unbounded, and efficiency obtains in the limit:*

$$\lim_{n\to\infty} \overline{v}^n = 1 - c.$$

In the appendix, we prove the proposition using the following steps. Based on a sand-

[24]In the literature on repeated games, this effect was noted by Matsushima (2001).

wich argument, we first show that, in the limit, $\hat{\rho}^n$ is of the deterministic form $\hat{r}^n = (1, \ldots, 1, 0, \ldots, 0)$. We next show that minimizing p_0 with respect to such deterministic cutoffs, only the two extreme constraints (IC_G^1) and (IC_G^n) need to be considered.[25] In a final step, we apply central limit arguments and Chebyshev's inequality to show that, in the limit as n goes to infinity, efficiency obtains for these deterministic cutoffs.

The Generic Case ($\lambda_L \in (0, 1)$). Proposition 7 implies that without a moral hazard problem, efficiency obtains when expanding the collective reputation over an infinite number of markets.

The next proposition shows that this is no longer true as soon as there is a slight moral hazard problem concerning the local player's effort. In this case, there is an upper bound on the extent of collective branding.

Proposition 8. *Suppose $\lambda_L > 0$. Then there is an upper bound $\bar{n} \in \mathbb{N}$ such that effort is implementable under collective branding of size n only if $n < \bar{n}$. Consequently, $\overline{V}^n > 0$ only if $n < \bar{n}$. Moreover, $\bar{n}$ is increasing in δ.*

Our final result returns to the main focus of our paper: identifying the inefficiencies of a collective reputation due to free-riding by local players. Proposition 7 shows that if collective reputation does not suffer from a free-riding problem, then, for any discount factor $\delta > c$, inefficiencies vanish as the collective brand size grows large. Our final proposition shows that if there is a free-riding problem ($\lambda_L > 0$), then, even as the discount factor δ approaches 1, the maximum per-period payoff $\bar{v}^n$ is bounded away from efficiency.

Proposition 9. *Suppose $\lambda_L > 0$. Then,*

$$\lim_{\delta \to 1} \bar{n} = \infty; \; \lim_{\delta \to 1} \hat{n} = \infty; \; \textit{and } \bar{v}^\infty \equiv \lim_{n \to \infty} \lim_{\delta \to 1} \bar{v}^n = \lambda_G(1 - c) + \lambda_L \bar{v}^I < 1 - c.$$

The proposition shows that, even in the limit, efficiency is not attained: the limiting payoff, $\bar{v}^\infty$, lies strictly below $1 - c$. Moreover, it reveals that we can decompose this limiting payoff by expressing it as a convex combination of the efficient payoff, $1 - c$, and the (inefficient) payoff under independent branding, $\bar{v}^I$, with the weights corresponding to the effort cost shares λ_G and λ_L.

This decomposability result indicates, that in the limit, it is as if we obtain the best of both worlds: implementing collective branding for the global player, yielding the efficient payoff $1 - c$, and, at the same time, independent branding for the local players, with its optimal but inefficient payoff of $\bar{v}^I$. The optimal calibration of the revenue shares is crucial

[25] Confirmed by simulations, we conjecture that this also holds with respect to an optimal non-deterministic $\hat{\rho}^n$ for finite n. The intractability of the Poisson's binomial distribution however prevents us from proving this analytically.

for this decomposability. This does, however, not mean that the two incentive problems are independent and do not interact; at all times incentive constraints of both types of players are binding.

6 Conclusion

We have developed a theory of collective brand reputation in a repeated game of imperfect public monitoring. A key novelty of our theory is the interaction between a global player (who takes costly actions to impact the quality of the entire product line) and local players (each of whom is able to affect the quality of only a single product). This makes the analysis applicable to a large set of economic environments in which such moral hazard problems are endemic, including platform markets, franchising, licensing, and team production.

While under independent branding, the quality signals relating to different products can be disentangled and are treated separately, they are effectively pooled under collective branding. If all of the effort costs are borne by the global player, the informativeness effect implies that collective branding is superior to independent branding. In that case, any inefficiency arising from imperfect monitoring vanishes in the limit as the collective brand size becomes large. By contrast, if all of the effort costs are borne by the local players, the free-rider effect implies that collective branding is never superior to independent branding, and is strictly inferior unless the discount factor is sufficiently large. In the generic case in which both types of players bear some of the effort costs, a careful calibration of revenues shares mediates the tradeoff between the beneficial informativeness effect and the harmful free-rider effect. Under optimal revenue sharing, collective branding is superior to independent branding as long as the share of the effort costs borne by the local players is sufficiently small or the discount factor sufficiently large. As the discount factor becomes large, the optimal size of the collective brand increases without bound. In the limit as both the discount factor and the collective brand size become large, the remaining inefficiency is equal to the local players' effort cost share multiplied by the inefficiency under independent branding. In that limit, it is thus as if the best of both worlds could be achieved: collective branding for the global player and independent branding for the local players.

Throughout the paper, we have assumed that the global player makes separate effort choices for each product/market. This may be a reasonable assumption for some applications (think of a franchisor's delivery of beef to hamburger outlets) but perhaps less so for some others (think of a headquarter's advertising in national media). If, under collective branding, the headquarter had to choose the same effort level in all markets, then only one aspect would need to be changed in our analysis: the global player would no longer have n incentive constraints but only a single one, namely (IC_G^{Cn}). It follows immediately that the resulting value under collective branding is weakly larger – and strictly larger for sufficiently large discount factors, than with separate effort decisions.[26] It is straightforward to show, however, that all of our propositions would continue to hold under that alternative assumption. In particular, the comparative statics and limiting values remain valid.

We have also assumed that effort choices are private information. An exciting avenue for future research consists in allowing for within-brand monitoring of effort choices. Such an analysis would, however, require a solution concept beyond PPE and therefore a more complex (and less well-understood) analytical framework. For instance, if the global player were to observe signals of local players' effort choices that are more informative than those observed by consumers, the global player's strategy would naturally depend on her private history (at least for the equilibrium to improve upon the PPE outcome in the absence of monitoring).

Because our analysis reveals the crucial role of revenue sharing, a further interesting question is to identify alternative modeling assumptions under which players benefit from more elaborate revenue sharing schemes than the ones we have analyzed. First, note that if the signals were fully contractible so that revenue shares could directly condition on them, then the players would be able to solve completely the moral hazard problem by using budget breakers, along the lines of Holmström (1982). Suppose instead, as we assume in the paper, that such direct conditioning on signal realizations is, due to their non-verifiability, infeasible. Then the players may, following the logic of relational contracting (e.g., Levin, 2003), try to exploit the repeated game structure to implement conditional revenue sharing implicitly through voluntary payments.[27] However, under our assumption that signals are fully uninformative about behavior of specific players (and, under collective branding, in specific markets), such relational contracts cannot help alleviate the moral hazard problem.[28] An open question though is whether this might be different if, under collective branding, all players could not only observe the aggregate signal but also attribute the individual signals to a specific local market. While such an assumption does not reflect our interpretation of collective branding, we expect relational contracts that implement voluntary payments from local markets with a bad signal to local markets with a good signal, to be sustainable and alleviate the moral hazard problem. However, for such outcomes to be attainable in a PPE, consumers must also be able to observe such voluntary transfers between producers in different markets, which seems unlikely to hold in practice.

Moreover, we assumed that effort choices of the global and local players are binary and perfectly complementary within a market. This assumption allowed us to focus on the

[26]Recall from our analysis that, in the best PPE, (IC_G^{C1}) is binding, and (IC_G^{Cn}) slack, for δ large.

[27]The relational contracting literature offers the insight that our focus on static revenue shares is without loss.

[28]Technically, the condition of "pairwise identifiability" (Fudenberg et al., 1994) fails in our context.

reputational team production problem across markets and abstract from the physical team production problem within a market. In a richer production structure with continuous effort choices that are imperfect complements, players must then, on the one hand, also solve the physical team production problem, but have, on the other hand, more punishments abilities to control reputational team production. We leave such an analysis for future research.

In our analysis, we have assumed that all markets are identical. Allowing for a market-specific signal structure (α_i, β_i) would enable us to study a number of novel questions such as *which* products to group together and sell under a common brand name, or how to optimally design aggregate brand-level quality signals. However, such an extension would have to deal with at least two analytical difficulties. First, this would require giving up the convenient restriction to symmetric equilibria. Second, this would require dealing with more complex Poisson's binomial distributions.[29]

Another interesting topic for future work consists in studying optimal task assignment within a collective brand. Suppose that production requires a continuum of tasks, indexed by $i \in [0, 1]$. Let $c_G(i)$ and $c_L(i)$ denote the effort cost for the global and local player, respectively, in performing task i. Suppose that $\Delta c(i) \equiv c_G(i) - c_L(i)$ is strictly decreasing in i, with $\Delta c(\hat{\imath}) = 0$ for some $\hat{\imath}$. First-best efficiency thus requires that tasks $[0, \hat{\imath})$ are performed by local players, and tasks $(\hat{\imath}, 1]$ by the global player. An implication of our analysis in Section 3 is that such "myopic" cost minimization is indeed optimal under independent branding. Under collective branding, however, our results imply that the global player should optimally take on more tasks as the joint value is decreasing in the share of the effort costs borne by the local players.[30] That is, optimal task assignment introduces a productive inefficiency under collective branding to mediate the free-rider problem.

[29]As Poisson's binomial distributions in general satisfy MLRP, we would expect the main insights of our analysis to carry over to such an extension.

[30]In the context of platform markets, the "Fulfillment by Amazon" (FBA) program may be understood through this lens. The FBA program, established in 2006, amounted to Amazon (as the global player) taking over the tasks of storage, shipping and handling returns from individual merchants (the local players).

Appendix A: Proofs

Proof of Proposition 1: Follows directly from the text. Q.E.D.

Proof of Lemma 1: First note that solving $\mathcal{P}^n$ disregarding all constraints, yields $\rho_s = 0$ for all s, but this violates all constraints. Because the optimization problem is linear, it follows that at least one of the constraints must be binding at an optimal solution.

Second, suppose to the contrary that ρ^C is optimal but such a $\bar{s}$ does not exist. Then there are $l < h$ such that $\rho_l < 1$ and $\rho_h > 0$, with ρ^C satisfying all the constraints of $\mathcal{P}^n$. Consider changing ρ^C to $\hat{\rho}^C$ by only lowering ρ_h by $\Delta\rho > 0$ and raising ρ_l by $\Delta\rho \cdot \mathbb{P}_n(s_h|0)/\mathbb{P}_n(s_l|0)$. This change does not affect $p_0 = \sum_s^n \mathbb{P}_n(s|0)\rho_s$. Therefore, the objective and the right-hand side of all the constraints remain unchanged. The left-hand side of the constraints change by

$$\left\{[\mathbb{P}_n(s_l|k) - \mathbb{P}_n(s_l|0)]\frac{\mathbb{P}_n(s_h|0)}{\mathbb{P}_n(s_l|0)} - [\mathbb{P}_n(s_h|k) - \mathbb{P}_n(s_h|0)]\right\}\frac{\Delta\rho}{k}.$$

After rewriting the term in curly brackets as

$$\left\{\frac{\mathbb{P}_n(s_l|k) - \mathbb{P}_n(s_l|0)}{\mathbb{P}_n(s_l|0)} - \frac{\mathbb{P}_n(s_h|k) - \mathbb{P}_n(s_h|0)}{\mathbb{P}_n(s_h|0)}\right\}\mathbb{P}_n(s_h|0) = \left\{\frac{\mathbb{P}_n(s_l|k)}{\mathbb{P}_n(s_l|0)} - \frac{\mathbb{P}_n(s_h|k)}{\mathbb{P}_n(s_h|0)}\right\}\mathbb{P}_n(s_h|0),$$

the MLRP of Poisson's binomial distribution implies that the term is strictly positive so that the left-hand side of each constraint strictly increases. As a result, $\hat{\rho}^C$ must also be optimal, since it attains the same objective value and all constraints are satisfied, even strictly so. The latter however contradicts the first observation that for any solution to $\mathcal{P}^n$ at least one constraint is binding.

To see the final claim of the lemma, note that given the cutoff structure, it follows that $p_k = \sum_{s=0}^n \mathbb{P}_n(s|k)\rho_s = \mathbb{P}_n(\tilde{s} < \bar{s}|k) + \mathbb{P}_n(\bar{s}|k)\rho_{\bar{s}} = (1-\rho_{\bar{s}})\mathbb{P}_n(\tilde{s} < \bar{s}|k) + \rho_{\bar{s}}\mathbb{P}_n(\tilde{s} \leq \bar{s}|k) \leq (1-\rho_{\bar{s}})\mathbb{P}_n(\tilde{s} < \bar{s}|k+1) + \rho_{\bar{s}}\mathbb{P}_n(\tilde{s} \leq \bar{s}|k+1) = \mathbb{P}_n(\tilde{s} < \bar{s}|k+1) + \mathbb{P}_n(\bar{s}|k+1)\rho_{\bar{s}} = p_{k+1}$ for all $k = 0, \ldots, n-1$, where the inequality follows from first-order stochastic dominance (i.e., equation (7)). Q.E.D.

Proof of Lemma 2: In order to derive the critical discount factor $\bar{\delta}^C$, we first compute for each of the constraints in (14) the minimum $\bar{\delta}_G^k$ such that there is an $\tilde{s}$ that fulfills it. In particular,

$$\bar{\delta}_G^k \equiv \frac{c}{\Delta_G^k n(\gamma_G - c) + c}, \text{ where } \Delta_G^k \equiv \max_{\tilde{s}}\left\{\sum_{s=0}^{\tilde{s}} \Delta_G(k,s)\right\}. \tag{19}$$

In order to see that $\gamma_G > \bar{\gamma}_G$ is a sufficient and necessary condition for $\Delta_G^k > 0$ for all

$k = 1, \ldots, n$, note that $\Delta_G(k, s) \geq 0$ only if $\Delta_G(k, 0) \geq 0$, since

$$\Delta_G(k,0) < 0 \;\Leftrightarrow \frac{\mathbb{P}_n(0|k)}{\mathbb{P}_n(0|0)} < 1 + \frac{kc}{n(\gamma_G - c)} \Rightarrow \frac{\mathbb{P}_n(s|k)}{\mathbb{P}_n(s|0)} < 1 + \frac{kc}{n(\gamma_G - c)} \Leftrightarrow \Delta_G(k,s) < 0,$$

where "$\Rightarrow$" follows from MLRP. Hence, $\Delta_G^k > 0$ if and only if $\Delta_G(k, 0) > 0$, where the latter is equivalent to

$$\gamma_G > c + \frac{kc}{n[((1-\beta)/\alpha)^k - 1]}.$$

Since the right hand side is decreasing in k,[31] a sufficient and necessary condition for $\Delta_G^k > 0$ for all $k = 1, \ldots, n$ is $\gamma_G > \bar{\gamma}_G$.

Likewise, compute for the constraint (15) the maximum $\bar{\delta}_L$ such that there is an $\tilde{s}$ that fulfills it. That is,

$$\bar{\delta}_L \equiv \frac{c}{\Delta_L(\gamma_L - c) + c}, \text{ where } \Delta_L \equiv \max_{\tilde{s}} \left\{ \sum_{s=0}^{\tilde{s}} \Delta_L(1,s) \right\} \tag{20}$$

and a sufficient and necessary condition for $\Delta_L > 0$ is $\gamma_L > \bar{\gamma}_L$.

It then follows that $\bar{\delta}^C$ is the maximum over all $\bar{\delta}_G^k$ and $\bar{\delta}_L$.

As shown, $\gamma_G > \bar{\gamma}_G$ and $\gamma_L > \bar{\gamma}_L$ imply $\Delta_G(k, 0) > 0$ for all $k = 1, \ldots, n$, and $\Delta_L(1, 0) > 0$. Noting that the right-hand sides of (14) and (15) vanish when δ approaches 1, the algorithm stops for $\tilde{s} = 0$, for δ close enough to 1, namely for $\delta \geq \bar{\delta}_0^C$, where

$$\bar{\delta}_0^C \equiv \max \left\{ \max_k \left\{ \frac{c}{c + n(\gamma_G - c)\Delta_G(k,0)} \right\}, \frac{c}{c + (\gamma_L - c)\Delta_L(1,0)} \right\}.$$

To see the last statement of the lemma, note that with $\bar{s} = 0$, (IC_G^{Ck}) reduces to

$$n\delta(\mathbb{P}_n(0|k) - \mathbb{P}_n(0|0))\rho_0 \frac{V(p_0, 1)}{k} \geq c.$$

Because $\mathbb{P}_n(0|k) = [(1-\beta)/\alpha]^k \alpha^n$ it follows from $(1-\beta)/\alpha > 1$ that $\mathbb{P}_n(0|k)$ is convex in k. As a result, the left hand side of (IC_G^{Ck}) is increasing in k while the right hand side is independent of k. Hence, if the constraint holds for $k = 1$, it holds for all $k > 1$. Q.E.D.

Proof of Proposition 2: The replication result derived in the body text shows that, for $\lambda_G = 1$ and $\rho_0^I \leq 1$, collective branding is weakly superior to independent branding as it can replicate the best PPE under independent branding. The fact that, except for the special case of $n = 2$ and $\rho_0^I = 1$, the superiority is strict, follows from Lemma 1, since

[31] Its derivative is of the same sign as $\psi(r, k) \equiv r^k - 1 - kr^k \log r$, where $r \equiv (1-\beta)/\alpha > 1$. As $\psi(1, k) = 0$ and $\partial\psi(r, k)/\partial r = -k^2 r^{k-1} \log r < 0$ for $r > 1$, it follows that $\psi(r, k) < 0$.

$\rho = (\rho_1(\rho_0^I), \ldots, \rho_n(\rho_0^I))$ violates the lemma's characterization of a solution to $\mathcal{P}^n$, except for the special case of $n = 2$ and $\rho_0^I = 1$. If $\rho_0^I > 1$ then $\overline{V}^I = 0$ so that a (weak) superiority of collective branding holds trivially since $\overline{V}^C \geq 0$. Whether the superiority is strict, depends on whether the discount factor δ is larger than the critical discount factor $\bar{\delta}^C$ in Lemma 2. In this respect, Proposition 2 implies that the critical discount factor $\bar{\delta}^C$ is smaller than the critical discount factor under independent branding, $\bar{\delta}^I$, and strictly so for any $n > 2$. Q.E.D.

Proof of Proposition 3: Follows directly from the text. Q.E.D.

Proof of Proposition 4: Follows directly from the text. Q.E.D.

Proof of Lemma 3: Suppose to the contrary that $\overline{V}^C > 0$ but $\hat{\gamma}_L \leq 1$, implying $\hat{\gamma}_G \geq 1$. Then, as argued, the program of minimizing p_0 w.r.t. ρ^C subject to all (IC_G^{Ck}) and (IC_L^C) has a solution $\rho^C = (1, \ldots, 1, \rho_{\bar{s}}, 0, \ldots, 0)$ such that (IC_L^C) binds, while all constraints (IC_G^{Ck}) are slack. Since $\partial V(p_0, \gamma_L)/\partial \gamma_L > 0$, raising γ_L slightly results in all constraints being slack, allowing to lower p_0 by reducing $\rho_{\bar{s}}$ (or if $\rho_{\bar{s}} = 0$ lowering $\bar{s}$). This contradicts that $\hat{\gamma}_L \leq 1$ is optimal. If $\hat{\gamma}_L \geq 1$ but no (IC_G^{Ck}) is binding, one can lower p_0 by the same procedure, whereas if (IC_L^C) is slack one can lower p_0 through a similar procedure by raising γ_G. We conclude that $\hat{\gamma}_G < 1 < \hat{\gamma}_L$ and are such that (IC_L^C) and at least one (IC_G^{Ck}) is binding.

Lemma 2 shows that for $\delta \in (\bar{\delta}_0^C, 1)$, the binding constraint must be (IC_G^{C1}). In this case, $(\hat{\gamma}_L, \hat{\gamma}_G)$ are such that both (IC_G^{C1}) and (IC_L^C) hold with equality, implying

$$\delta(p_1 - p_0)V(p_0, \hat{\gamma}_G)n = c = \delta(p_1 - p_0)V(p_0, \hat{\gamma}_L) \;\Leftrightarrow\; n(\hat{\gamma}_G - c) = \hat{\gamma}_L - c.$$

Combining the latter equation with the identity $\gamma_L = (1 - \lambda_G \gamma_G)/(1 - \lambda_G)$ yields $\hat{\gamma}_G = \tilde{\gamma}_G$ and $\hat{\gamma}_L = \tilde{\gamma}_L$. Note that for $\delta < \bar{\delta}_0^C$ the constraint (IC_G^{C1}) may be slack at the optimum. Hence, in general, we have

$$\delta(p_1 - p_0)V(p_0, \hat{\gamma}_G)n \geq c = \delta(p_1 - p_0)V(p_0, \hat{\gamma}_L) \;\Leftrightarrow\; n(\hat{\gamma}_G - c) \geq \hat{\gamma}_L - c.$$

Combining this latter inequality with our previous finding that $\hat{\gamma}_G < 1 < \hat{\gamma}_L$ yields the result $\hat{\gamma}_G \in [\tilde{\gamma}_G, 1)$ and $\hat{\gamma}_L \in (1, \tilde{\gamma}_L]$. Q.E.D.

Proof of Proposition 5: By Proposition 3, collective branding with proportional rewards can attain the value $\overline{V}^I$, whenever $\rho_0^I \leq \alpha^{n-1}$. Lemma 3 then implies that $\overline{V}^C > \overline{V}^I$, since it shows that proportional rewards are strictly suboptimal, meaning that collective branding can attain a value exceeding $\overline{V}^I$. The result then follows from the observation that $\rho_0^I \leq \alpha^{n-1}$

if and only if $\delta \geq \bar{\delta}$, where

$$\bar{\delta} = \frac{c}{c + \alpha^{n-1}(1 - \alpha - \beta - (1 - \beta)c)}.$$

Q.E.D.

Proof of Proposition 6: Suppose the model's parameters are such that $\overline{V}^C > 0$, implying that there is an optimal triple $(\hat{\gamma}_G, \hat{\gamma}_L, \hat{\rho})$ to $\mathcal{P}^C$. In particular, for this triple $(\hat{\gamma}_G, \hat{\gamma}_L, \hat{\rho})$, the constraint $(IC_L^C(\hat{\gamma}_L))$ as well as at least one constraint $(IC_G^{Ck}(\hat{\gamma}_G))$ are binding. Now consider an increase in the parameter λ_G, while keeping γ_G constant at $\hat{\gamma}_G$. These changes imply a decrease in λ_L together with an increase in π_G. From $\gamma_L = (1 - \lambda_G \gamma_G)/(1 - \lambda_G)$, it however follows that $\partial \gamma_L / \partial \lambda_G > 0$ so that the overall effect on γ_L is positive. Hence, (IC_L^C) is relaxed, implying that we can, in fact, also increase γ_G slightly above $\hat{\gamma}_G$, thereby relaxing all constraints so that $\hat{\rho}$ together with the raised γ_G and γ_L lead to the same p_0 but with all constraints satisfied with strict inequality. By Lemma 3, the triple is suboptimal, meaning there is a different triple leading to a strictly lower p_0, implying a strictly larger $\overline{V}^C$. Q.E.D.

Proof of Proposition 7: Let $r^s \in \mathbb{R}^{n+1}$ denote deterministic cutoffs ρ of the form $(1, \ldots, 1, 0, \ldots, 0)$. That is, r^s is an $(n+1)$-dimensional vector with the first s entries being 1 and the remaining $n + 1 - s$ entries being 0. Denote by $R^n = \{r^0, r^1, \ldots, r^{n+1}\} \subset \mathbb{R}^{n+1}$ the set of all r^s for a given n.

Given n and $\gamma_G^n = 1$, recall that we have

$$\hat{\rho}^n = \arg\min_{\rho} p_0(\rho) \text{ s.t. } (IC_G^{n;1}(1)), \ldots, (IC_G^{n;n}(1)).$$

In addition to this minimization problem, consider for a given n and $\gamma_G^n = 1$ the problems

$$\tilde{\rho}^n = \arg\min_{\rho} p_0(\rho) \text{ s.t. } (IC_G^{n;1}(1)) \text{ and } (IC_G^{n;n}(1));$$

$$\hat{r}^n = \arg\min_{r^s \in R^n} p_0(r^s) \text{ s.t. } (IC_G^{n;1}(1)), \ldots, (IC_G^{n;n}(1));$$

$$\tilde{r}^n = \arg\min_{r^s \in R^n} p_0(r^s) \text{ s.t. } (IC_G^{n;1}(1)) \text{ and } (IC_G^{n;n}(1)).$$

Because the first problem is less stringent than the minimization problem underlying $\hat{\rho}^n$, whereas the second is more stringent, it follows

$$p_0(\tilde{\rho}^n) \leq p_0(\hat{\rho}^n) \leq p_0(\hat{r}_n).$$

Moreover, if $\tilde{\rho}^n$ has the cutoff at $\bar{s}$, then $\tilde{r}^n = r^{\bar{s}+1}$ so that

$$p_0(\tilde{r}^n) - p_0(\tilde{\rho}^n) = \mathbb{P}_n\{\bar{s}|0\}(1 - \tilde{\rho}^n_{\bar{s}}).$$

Since $\lim_{n\to\infty} \mathbb{P}_n\{\bar{s}|0\} = 0$, it follows

$$\lim_{n\to\infty} p_0(\tilde{\rho}^n) = \lim_{n\to\infty} p_0(\tilde{r}^n). \tag{21}$$

Because the minimization problem associated with $\tilde{r}^n$ is a relaxed version of the minimization problem associated with $\hat{r}_n$, it holds $p_0(\hat{r}_n) \geq p_0(\tilde{r}^n)$. The next lemma shows that, in fact, $p_0(\hat{r}_n) = p_0(\tilde{r}^n)$.

Lemma 4. *The minimizer $\hat{r}^n$ minimizes p_0 subject to only the incentive constraints $(IC_G^{n;1}(1))$ and $(IC_G^{n;n}(1))$.*

Proof. The statement is trivially satisfied for $n = 2$, so suppose $n > 2$. Rewriting $(IC_G^{n;k}(1))$ as

$$\frac{p_k^n - p_0^n}{k} \geq \frac{c}{\delta n V(p_0, 1)} \tag{22}$$

shows that if p_k is convex in k, in the sense that the incremental differences

$$p_{k+2}^n - p_{k+1}^n - (p_{k+1}^n - p_k^n) \tag{23}$$

are positive, then the left hand side of (22) is increasing in k so that $(IC_G^{n;1}(1))$ implies all other constraints and, hence, at most $(IC_G^{n;1}(1))$ can be binding. If, in contrast, p_k is concave in k, then the left hand side is decreasing in k so that $(IC_G^{n;n}(1))$ implies all other constraints and, hence, at most $(IC_G^{n;n}(1))$ can be binding.

We next argue that, for any $r^{\bar{s}} \in R^n$, the curvature of $p_k(r^{\bar{s}})$ w.r.t. k, i.e., the sign of (23), depends on the cutoff $\bar{s}$. To see this, note first that the recursive structure (6) of $\mathbb{P}_n(s|k)$ implies that the incremental differences (23) rewrites as:

$$\begin{aligned} p_{k+2}^n - p_{k+1}^n - (p_{k+1}^n - p_k^n) =& \sum_{s=0}^{\bar{s}} [\mathbb{P}_n(s|k+2) + \mathbb{P}_n(s|k) - 2\mathbb{P}_n(s|k+1)] \\ =&(1-\alpha-\beta)^2\{[\mathbb{P}_{n-2}(\bar{s}|k) - \mathbb{P}_{n-2}(\bar{s}-1|k)].\} \end{aligned} \tag{24}$$

The single peakedness of $\mathbb{P}_{n-2}(.|k)$ implies that $\mathbb{P}_{n-2}(s|k)$ is increasing in s for all s smaller than the mode, m_k, of $\mathbb{P}_{n-2}(.|k)$ and decreasing in s for all s larger than m_k. Since m_k is decreasing in k, m_k lies in between the modes of the binomial distributions $B(n-2, \beta)$ and $B(n-2, (1-\alpha))$, i.e., $m_k \in [m_0, m_{n-2}]$ with $m_0 = \lfloor (n-1)\beta \rfloor$ and $m_{n-2} = \lfloor (n-1)(1-\alpha) \rfloor$, where $\lfloor x \rfloor$ is the greatest integer less than or equal to x. Therefore, if $\bar{s} \leq m_0$, then (24) is positive for all k and, hence, the incremental differences are positive for all k, implying p_k is convex. Similarly, if $\bar{s}-1 > m_{n-2}$, then (24) is negative for all k and, hence, the incremental differences are negative for all k, implying p_k is concave. For $\bar{s} \in [m_0, m_{n-2}+1]$, the sign of (24) may depend on k. The single peakedness of $\mathbb{P}_{n-2}(.|k)$ together with the decreasing mode m_k implies however that (24) can switch sign at most once as k increases and only from positive to negative. If so, there is a $\bar{k}$ such that p_k is convex for all $k \leq \bar{k}$ and concave for all $k \geq \bar{k}$. This means that for $k \leq \bar{k}$, ($IC_G^{n;1}(1)$) implies ($IC_G^{n;k}(1)$), and for $k \geq \bar{k}$, ($IC_G^{n;n}(1)$) implies ($IC_G^{n;k}(1)$). □

A direct corollary of Lemma 4 is that

$$p_0^n(\tilde{\rho}^n) \leq p_0^n(\hat{\rho}^n) \leq p_0^n(\hat{r}_n) = p_0^n(\tilde{r}^n).$$

Since this string of inequalities holds for all n, it holds also in the limit, implying that

$$\lim_{n\to\infty} p_0^n(\tilde{\rho}^n) \leq \lim_{n\to\infty} p_0^n(\hat{\rho}^n) \leq \lim_{n\to\infty} p_0^n(\hat{r}_n) = \lim_{n\to\infty} p_0^n(\tilde{r}^n).$$

By (21) and a sandwich theorem, it then follows that

$$\lim_{n\to\infty} p_0^n(\tilde{\rho}^n) = \lim_{n\to\infty} p_0^n(\hat{\rho}^n) = \lim_{n\to\infty} p_0^n(\hat{r}_n) = \lim_{n\to\infty} p_0^n(\tilde{r}^n).$$

The next lemma shows that, for large enough n, we can pick deterministic cutoffs such that p_0 is arbitrarily close to zero, while the constraints are all slack, provided that $\delta > c$.

Lemma 5. *Suppose $\delta > c$ so that $\bar{\varepsilon} \equiv (\delta - c)/\delta > 0$. Then, for all $\varepsilon \in (0, \bar{\varepsilon})$*

$$\exists \bar{n} \in \mathbb{N} : \forall n > \bar{n},\ \exists \bar{s}(n) \leq n : p_0^n(r^{\bar{s}(n)}) < \varepsilon \text{ and } (IC_G^{n;1}(1)) \text{ and } (IC_G^{n;n}(1)) \text{ are slack.}$$

Proof. Fix $\varepsilon > 0$ and let

$$K = \frac{c}{\delta V(\varepsilon, 1)(1-\alpha-\beta)}.$$

Denote the pdf of the normal distribution $\mathcal{N}(\mu, \sigma^2)$ with mean $\mu = n(1-\alpha)$ and variance $\sigma^2 = n\alpha(1-\alpha)$ by

$$\varphi_n(s) = \frac{1}{\sqrt{2\pi n\alpha(1-\alpha)}} e^{-\frac{1}{2}\frac{(s-n(1-\alpha))^2}{n\alpha(1-\alpha)}}.$$

Then there is an $\tilde{n}_1$ such that for all $n > \tilde{n}_1$, the equation

$$\varphi_n(\tilde{s}(n)) = \frac{K}{(1-\varepsilon)n}$$

has the solution $\tilde{s}(n) = \mu - \sigma h(n) < (1-\alpha)n - 1$ with

$$h(n) \equiv \sqrt{\ln\left(\frac{n}{2\pi\alpha(1-\alpha)K^2/(1-\varepsilon)^2}\right)}.$$

By noting that $\lim_{n\to\infty} h(n) = \infty$, Chebyshev's inequality implies that there is an $\tilde{n}_2(> \tilde{n}_1)$ such that for all $n > \tilde{n}_2$ it holds

$$\Phi_n(\tilde{s}(n)) \equiv \Pr\{s \leq \tilde{s}(n)\} = \int_{-\infty}^{\tilde{s}(n)} \varphi_n(s)ds < \varepsilon/2.$$

Since the binomial distribution $B(n,(1-\alpha))$ converges in distribution to $\mathcal{N}(\mu, \sigma^2)$, there is an $\tilde{n}_3(> \tilde{n}_2)$ such that for any $n > \tilde{n}_3$

$$|\mathbb{P}_n\{s \leq \bar{s}(n)|0\} - \Phi_n(\tilde{s}(n))| < \varepsilon/2,$$

where $\bar{s}(n) = \lceil \tilde{s}(n) \rceil$ is the smallest integer greater than $\tilde{s}(n)$. Hence, for any $n > \tilde{n}_3$

$$p_0^n(r^{\bar{s}(n)}) = \mathbb{P}_n\{s \leq \bar{s}(n)|0\} < \Phi_n(\tilde{s}(n)) + \varepsilon/2 < \varepsilon.$$

It remains to be shown that $r^{\bar{s}(n)}$ satisfies $(IC_G^{n;1}(1))$ and $(IC_G^{n;n}(1))$ with slackness.

The constraint $(IC_G^{n;1}(1))$ for $r^{\bar{s}(n)}$ rewrites as

$$\mathbb{P}_{n-1}(\bar{s}(n)|0) \geq \frac{1}{n}\frac{c}{\delta V(p_0(r^{\bar{s}(n)}), 1)(1-\alpha-\beta)}.$$

Since $V(p_0^n, 1)$ is strictly decreasing in p_0^n and $p_0^n(r^{\bar{s}(n)}) < \varepsilon$, the inequality holds strictly if

$$\mathbb{P}_{n-1}(\bar{s}(n)|0) \geq K/n.$$

The de Moivre-Laplace theorem implies there is an $\tilde{n}_4(> \tilde{n}_3)$ such that for all $n > \tilde{n}_4$

$$1 - \varepsilon < \frac{\mathbb{P}_{n-1}(\bar{s}(n)|0)}{\varphi_n(\bar{s}(n))} < 1 + \varepsilon.$$

Hence, for $n > \tilde{n}_4$ it holds

$$\mathbb{P}_{n-1}(\bar{s}(n)|0) > (1-\varepsilon)\varphi_n(\bar{s}(n)) \geq (1-\varepsilon)\varphi_n(\tilde{s}(n)) = K/n,$$

where the second inequality follows since, $\bar{s}(n) \leq \lceil \tilde{s}(n) \rceil \leq \lceil (1-\alpha)n - 1 \rceil \leq (1-\alpha)n$. This shows that $r^{\bar{s}(n)}$ satisfies $(IC_G^{n;1})$ with slackness for all $n > \tilde{n}_4$.

Recall $(IC_G^{n;n})$ for $r^{\bar{s}(n)}$ rewrites as

$$p_n^n(r^{\bar{s}(n)}) - p_0^n(r^{\bar{s}(n)}) \geq \frac{c}{\delta V(p_0(r^{\bar{s}(n)}), 1)}.$$

Since $\lim_{n\to\infty} \tilde{s}(n)/n = (1-\alpha) > \beta$, it follows by the law of large numbers that for any $\tilde{\varepsilon} > 0$, there is an $\tilde{n}_5(> \tilde{n}_4)$ so that for all $n > \tilde{n}_5$, we have

$$p_n^n(r^{\bar{s}(n)}) = \sum_{s=0}^{\bar{s}(n)} \mathbb{P}_n(s|n) > 1 - \tilde{\varepsilon}.$$

In particular, for $\tilde{\varepsilon} \in (0, (\delta - c - \delta\varepsilon)/(\delta(1-c)))$, it follows

$$p_n^n(r^{\bar{s}(n)}) - p_0^n(r^{\bar{s}(n)}) > 1 - \tilde{\varepsilon} - \varepsilon > \frac{c(1-\delta+\delta\varepsilon)}{\delta(1-c)} = \frac{c}{\delta V(\varepsilon, 1)} > \frac{c}{\delta V(p_0^n(r^{\bar{s}(n)}), 1)}.$$

This confirms that $(IC_G^{n;n}(1))$ for $r^{\bar{s}(n)}$ holds with slackness for all $n > \tilde{n}_5$. Taking $\bar{n} = \tilde{n}_5$ completes the lemma. □

Noting that p_0^n arbitrarily close to 0 means that the average per period payoffs, v^n, is arbitrarily close to $1-c$, then yields the Proposition. Q.E.D.

Proof of Proposition 8: If $\lambda_L > 0$ and $\overline{V}^n > 0$, then there is a triple $(\hat{\gamma}_G^n, \hat{\gamma}_L^n, \hat{\rho}^n)$ where $\hat{\rho}^n = (1, \ldots, 1, \rho_{\bar{s}^n}, 0, \ldots, 0) \in \mathbb{R}^{n+1}$ with $\bar{s}^n \in \{0, \ldots, n-1\}$ such that incentive constraint $(IC_L^C(\hat{\gamma}_L^n))$ is satisfied. That is,

$$p_1^n(\hat{\rho}^n) - p_0^n(\hat{\rho}^n) \geq \frac{c}{\delta V(p_0^n(\hat{\rho}^n), \hat{\gamma}_L^n)}. \tag{25}$$

Because $p_0^n(\hat{\rho}^n) \geq 0$ and $\hat{\gamma}_L^n \leq 1/\lambda_L$, we have

$$V(p_0^n(\hat{\rho}^n), \hat{\gamma}_L^n) = \frac{\hat{\gamma}_L^n - c}{1 - \delta(1 - p_0^n(\hat{\rho}^n))} \leq \frac{1/\lambda_L - c}{1-\delta},$$

so that the right-hand side of (25) is larger than some lower bound that is strictly larger than zero and independent of n.

We next argue that the left-hand side, $p_1^n - p_0^n$, of (25) goes to zero as n grows arbitrarily large. To see this, apply the recursive structure (6) to obtain

$$p_1^n(\hat{\rho}^n) - p_0^n(\hat{\rho}^n) = (1-\alpha-\beta)[\mathbb{P}_{n-1}(\bar{s}^n - 1|0)(1 - \hat{\rho}_{\bar{s}^n}) + \mathbb{P}_{n-1}(\bar{s}^n|0)\hat{\rho}_{\bar{s}^n}]. \tag{26}$$

Since $\mathbb{P}_{n-1}(.|0)$ is a binomial distribution of $n-1$ trials with success probability $1-\alpha$, the individual probability $\mathbb{P}_{n-1}(s|0)$ goes to zero for any s as n grows arbitrarily large. Equation

(26) then implies that $p_1^n(\hat{\rho}^n) - p_0^n(\hat{\rho}^n)$ goes to zero as n grows arbitrarily large, which implies that there exists an $\hat{n}$ such that (25) is violated for all $n > \hat{n}$. Consequently, there is an upper bound on n such that $\mathcal{P}^n$ has a solution. Since n is an integer, there is actually a lower upper bound $\bar{n}$ such that $\mathcal{P}^n$ has a solution if and only if $n < \bar{n}$. Note that if for parameters $(\alpha, \beta, c, \delta, \lambda_L, n)$ the program $\mathcal{P}^n$ has a solution then it also has a solution when δ increases, implying that $\bar{n}$ is increasing in δ. Q.E.D.

Proof of Proposition 9: We first prove the second statement (which implies the first statement) by showing that for any n, there is a $\bar{\delta}(n)$ such that for all δ that exceed $\bar{\delta}(n)$, it holds $\hat{n} > n$. To see this, fix n and take $\gamma_L = \tilde{\gamma}_L^n$. Lemma 2 implies that for $\delta > \bar{\delta}_0^n$, we have: $\hat{\rho}_0^n < 1$, $\hat{\rho}_i^n = 0$ for all $i = 1, \ldots, n$, $(IC_G^{n;1}(\tilde{\gamma}_G^n))$ holds with equality while all other $(IC_G^{n;k}(\tilde{\gamma}_G^n))$ are slack. Moreover, by Lemma 3, $\hat{\gamma}_L = \tilde{\gamma}_L^n$ so that $(IC_L^n(\tilde{\gamma}_L^n))$ holds with equality for $\hat{\rho}^n$. Using this latter equality yields that for any $\delta > \bar{\delta}_0^n$ we have

$$\hat{\rho}_0^n = \frac{c/\delta - c}{(\mathbb{P}_n(0|1) - \mathbb{P}_n(0|0))(\tilde{\gamma}_L^n - c) - \mathbb{P}_n(0|0)c} = \frac{c/\delta - c}{\alpha^{n-1}[(1-\alpha-\beta)(\tilde{\gamma}_L^n - c) - \alpha c]}. \tag{27}$$

Note that the right-hand side converges to zero as δ goes to one. Hence, given n, we can find a $\bar{\delta}(n) < 1$ so that for all $\delta > \bar{\delta}(n)$, we have $\hat{\rho}_0^n < \alpha$.

We next argue that for all $\delta > \bar{\delta}(n)$, we must have $n \neq \hat{n}$, because, already for brand size $n+1$, we can, given δ, find a ρ^{n+1} that yields a strictly lower p_0^{n+1}. To see this, consider $\rho^{n+1} = (\hat{\rho}_0^n/\alpha, 0, \ldots, 0)$ and $\gamma_L = \tilde{\gamma}_L^{n+1}$. It follows that

$$p_0^{n+1}(\rho^{n+1}) = \mathbb{P}_{n+1}(0|0)\hat{\rho}_0^n/\alpha = \alpha^{n+1}\hat{\rho}_0^n/\alpha = \alpha^n\hat{\rho}_0^n = \mathbb{P}_n(0|0)\hat{\rho}_0^n = p_0^n(\hat{\rho}^n)$$

and

$$p_1^{n+1}(\rho^{n+1}) = \mathbb{P}_{n+1}(0|1)\hat{\rho}_0^n/\alpha = \alpha^n(1-\beta)\hat{\rho}_0^n/\alpha = \alpha^{n-1}(1-\beta)\hat{\rho}_0^n = \mathbb{P}_n(0|1)\hat{\rho}_0^n = p_1^n(\hat{\rho}^n).$$

Hence,

$$\begin{aligned}[p_1^{n+1}(\rho^{n+1}) - p_0^{n+1}(\rho^{n+1})]\tilde{\gamma}_L^{n+1} - p_0^{n+1}(\rho^{n+1})c &= [p_1^n(\hat{\rho}^n) - p_0^n(\hat{\rho}^n)]\tilde{\gamma}_L^{n+1} - p_0^n(\hat{\rho}^n)c \\ &> [p_1^n(\hat{\rho}^n) - p_0^n(\hat{\rho}^n)]\tilde{\gamma}_L^n - p_0^n(\hat{\rho}^n)c = c/\delta - c,\end{aligned}$$

where the inequality holds because $\tilde{\gamma}_L^{n+1} > \tilde{\gamma}_L^n$. This implies that $(IC_L^{n+1}(\tilde{\gamma}_L^{n+1}))$ is slack. As $\gamma_L = \tilde{\gamma}_L^{n+1}$, the left-hand sides of $(IC_L^{n+1}(\tilde{\gamma}_L^{n+1}))$ and $(IC_G^{n+1;1}(\tilde{\gamma}_G^n))$ are equal to each other, implying that $(IC_G^{n+1;1}(\tilde{\gamma}_G^n))$ and, thus, all the other $(IC_G^{n+1;k}(\tilde{\gamma}_G^n))$ are slack as well. For the optimal $\hat{\rho}^{n+1}$ it therefore holds $p_0^{n+1}(\hat{\rho}^{n+1}) < p_0^n(\hat{\rho}^n)$. For $\delta > \bar{\delta}(n)$, brand size $n+1$ is therefore superior to brand size n. Since this argument holds for any n, it follows that $\hat{n}$ cannot be finite and $\lim_{\delta \to 1} \hat{n} = \infty$.

To see the final statement of the proposition, first fix n and $\gamma_L = \tilde{\gamma}_L^n$, and pick a δ arbitrarily close to 1, then, as discussed above, $\hat{\rho}^n$ is such that $\hat{\rho}_0^n$ equals (27) and $\hat{\rho}_i^n = 0$ for all $i = 1, \ldots, n$. The average per-period value can therefore be written as

$$\overline{v}^n = \frac{(1-\delta)(1-c)}{1-\delta+\delta\alpha^n\hat{\rho}_0^n} = \frac{(1-c)}{1+\frac{c\alpha}{(1-\alpha-\beta)(\tilde{\gamma}_L^n-c)-\alpha c}} = (1-c) - \frac{\alpha c(1-\lambda_L+\lambda_L n)}{(1-\alpha-\beta)n}.$$

Hence,

$$\lim_{n\to\infty} \overline{v}^n = 1 - c - \lambda_L \frac{\alpha c}{1-\alpha-\beta}.$$

Q.E.D.

Appendix B

In this appendix, we apply the abstract methods of decomposability and self-generation developed in Abreu, Pearce, and Stacchetti (1990) for the case of two markets, $n = 2$, under collective branding. In particular, we focus on the implementation of equilibrium outcomes in which players choose the cooperative actions $b_1 = e_{G,1} = e_{L,1} = b_2 = e_{G,2} = e_{L,2} = 1$ in the first period of the repeated game, both for symmetric and asymmetric perfect public equilibrium outcomes.

For fixed revenue shares $\pi_G = 1 - \pi_L$, the infinitely repeated game with imperfect monitoring has 5 players – the two local players, the global player, and the two consumers. We will refer to local player 1 as player 1, local player 2 as player 2, the global player as player 3, consumer 1 as player 4, and consumer 2 as player 5. Players 1, 2, and 3 are long-lived players, whereas the consumers, player 4 and 5, are short-lived. Except for the global player, all players have a binary action set, $A_1 = A_2 = A_4 = A_5 = \{0, 1\}$. The global player, as player 3, has an action set containing four actions that we can express as binary numbers, denoting in which of the two markets the global player picks effort: i.e., $A_3 = \{e_{G1}e_{G2}\}_{e_{G1},e_{G2}\in\{0,1\}} = \{00, 10, 01, 11\}$. Expressing player 3's action as a binary number, we can represent a pure action profile a as an element from $\{0, 1\}^6$ and the set of pure action profiles contains $2^6 = 64$ elements.

Because in equilibrium the short-lived players play myopic best replies, the set of *feasible* pure action profiles in the stage game of the overall repeated game is smaller. As explained in the main text, a consumer in market i buys if and only if the local and global player exert effort in market i. As a result, the set of feasible pure action profiles, $\mathbf{B}$, contains $2^4 = 16$ elements and for any feasible pure action profile consumers obtain a payoff of zero. Restricting attention to the set of feasible pure action profiles allows us to focus on the long-lived players, 1, 2, and 3, while ensuring equilibrium behavior of the short-run players. Concerning the long-lived players, the feasible action profile $a = (e_{L1}, e_{L2}, e_{G1}, e_{G2}, b_1, b_2) \in \mathbf{B}$ yields the following stage payoffs to the three (long-lived) players:

$$u_1(a) = a_{C1}\pi_L - \lambda_L e_{L1}c;\; u_2(a) = a_{C2}\pi_L - \lambda_L e_{L2}c;\; u_3(a) = (b_1 + b_2)\pi_G - \lambda_G(e_{G1} + e_{G2})c.$$

Note that, restricted to the feasible pure action profiles in $\mathbf{B}$, each player can guarantee himself at least a zero payoff by not exerting any effort. Moreover, by not exerting any effort, any pair of players can ensure that the other player gets at most a zero payoff. Hence, the minmax-payoff of each player is zero.

The observable signals are the aggregated quality reports $s = s_1 + s_2 \in \{0, 1, 2\}$ and the uniformly distributed public correlation device $r \in [0, 1]$. Given the action profile a, the perfect complementarity of efforts imply that the probability of signal $s_i \in \{0, 1\}$ in market

i is

$$\mathbb{P}\{s_i = 1|a\} = (1-\alpha)a^i a^{i+2} + \beta(1 - a^i a^{i+2}),$$

where a^i is the $i-$th element of the action profile $a = (e_{L1}, e_{L2}, e_{G1}, e_{G2}, b_1, b_2) \in \mathbf{B}$. Following Mailath and Samuelson (2006,p.253), we combine the signal s and r into one continuous signal y by defining

$$y = s_1 + s_2 + r \in Y \equiv [0, 3].$$

We denote the density of this continuous signal over the support Y by $\rho(y|a)$. Since the distribution of r is uniform, the density $\rho(y|a)$ is the step function:

$$\rho(y|a) = \begin{cases} \mathbb{P}\{s = 0|a\} & \text{, if } y \in [0, 1) \\ \mathbb{P}\{s = 1|a\} & \text{, if } y \in [1, 2) \\ \mathbb{P}\{s = 2|a\} & \text{, if } y \in [2, 3]. \end{cases}$$

Following Abreu, Pearce, and Stacchetti (1990), we define an action profile $a \in \mathbf{B}$ as enforceable on $\mathcal{W} \subset \mathbb{R}^3_+$ if there exists a (Lebesgue measurable) mapping $\gamma : Y \to \mathcal{W}$ such that for any $i = 1, 2, 3$,

$$\begin{aligned} V_i(a, \gamma) &\equiv (1-\delta)u_i(a) + \delta \int_0^3 \gamma_i(y)\rho(y|a)dy \\ &\geq (1-\delta)u_i(a_i', a_{-i}) + \delta \int_0^3 \gamma_i(y)\rho(y|a_i', a_{-i})dy \text{ for all } a_i' \in A_i. \end{aligned}$$

Note that, given the density $\rho(y|a)$, it holds for any $a \in \mathbf{B}$ which is enforceable on $\mathcal{W}$ that

$$\int_0^3 \gamma_i(y)\rho(y|a)dy = \sum_{j=0}^{2} \int_j^{j+1} \gamma_i(y)\rho(y|a)dy = \sum_{j=0}^{2} \mathbb{P}\{s = j|a\} \int_j^{j+1} \gamma_i(y)dy = \sum_{j=0}^{2} \mathbb{P}\{s = j|a\} w_j,$$

where $w_j = \int_j^{j+1} \gamma_i(y)dy$ lies in the convex hull of $\mathcal{W}$. Consequently, the definition of enforceability in our framework is equivalent to saying that an action profile $a \in \mathbf{B}$ is enforceable on $\mathcal{W} \in \mathbb{R}^3_+$ if there exists a triple w_1, w_2, w_3 in the convex hull of $\mathcal{W}$ such that

$$\begin{aligned} V_i(a, \gamma) &\equiv (1-\delta)u_i(a) + \delta \sum_{j=0}^{2} \mathbb{P}\{s = j|a\} w_j \\ &\geq (1-\delta)u_i(a_i', a_{-i}) + \delta \sum_{j=0}^{2} \mathbb{P}\{s = j|a_i', a_{-i}\} w_j \text{ for all } a_i' \in A_i. \end{aligned} \tag{28}$$

The perfect complementarity in the effort levels implies that any action profile for which, in some market, effort is only supplied by one player is not enforceable. As a result, only the following 4 action profiles of the 16 feasible action profile in **B** are enforceable:

$$(0,0,00,0,0);(1,0,10,1,0);(0,1,01,0,1);(1,1,11,1,1);$$

with associated enforceable payoff vectors

$$\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix};\begin{pmatrix} u_L \\ 0 \\ u_G \end{pmatrix};\begin{pmatrix} 0 \\ u_L \\ u_G \end{pmatrix};\begin{pmatrix} u_L \\ u_L \\ 2u_G \end{pmatrix},$$

where $u_L = \pi_L - \lambda_L c$ and $u_G = \pi_G - \lambda_G c$.

Abreu, Pearce, and Stacchetti (1990) show that the set of equilibrium payoffs, $\mathcal{E}(\delta)$, is a subset of the convex hull of the enforceable payoff vectors. Because these payoff vectors lie on a two dimensional plane within $\mathbb{R}^3$, we can express any point $w \in \mathbb{R}^3$ in the convex hull of these four points by a unique pair of scalars $(\mu_1, \mu_2) \in [0,1]^2$ such that

$$w = \hat{w}(\mu_1, \mu_2) \equiv \mu_1 \begin{pmatrix} u_L \\ 0 \\ u_G \end{pmatrix} + \mu_2 \begin{pmatrix} 0 \\ u_L \\ u_G \end{pmatrix}.$$

Hence, the convex hull is the set $\hat{W} \equiv \{\hat{w}(\mu_1, \mu_2) | \mu_1, \mu_2 \in [0,1]\}$ and the aggregate payoff associated with any point $\hat{w}(\mu_1, \mu_2)$ is $V(\mu_1, \mu_2) = (\mu_1 + \mu_2)(u_L + u_G)$. As a result, the equilibrium payoff that maximizes aggregate payoffs is attained by a solution (μ_1^*, μ_2^*) of the following program:

$$\mathcal{P} : \max_{\mu_1, \mu_2} V(\mu_1, \mu_2) \text{ s.t. } \hat{w}(\mu_1, \mu_2) \in \mathcal{E}(\delta).$$

Apart from the fact that $\mathcal{E}(\delta) \subset \hat{W}$, the presence of the correlation device r implies that the set $\mathcal{E}(\delta)$ is convex. Moreover, since the two markets are symmetric, $\mathcal{E}(\delta)$ exhibits the symmetry that $w(\mu_1, \mu_2) \in \mathcal{E}(\delta)$ implies $w(\mu_2, \mu_1) \in \mathcal{E}(\delta)$. From the convexity and this symmetry of $\mathcal{E}(\delta)$, it follows that there is a symmetric solution $\mu_1^* = \mu_2^* = \mu^*$ to problem $\mathcal{P}$ with an associated payoff vector w^* that is symmetric in the sense that $w^* = \hat{w}(\mu, \mu)$ for some $\mu \in [0,1]$. Note also that for any $w = (w_1, w_2, w_3) \in \mathcal{E}(\delta)$ we have $w = \hat{w}(w_1/u_L, w_2/u_L)$ so that be defining

$$M(\delta) \equiv \{(w_1/u_L, w_2/u_L) | (w_1, w_2, w_3) \in \mathcal{E}(\delta)\},$$

we obtain an equivalent representation of $\mathcal{E}(\delta)$ in terms of pairs (μ_1, μ_2).

Implementation of $a^{11} \equiv (1,1,1,1,1,1)$ in the first period. We next study an equilibrium in which a^{11} is the aggregate-payoff-maximizing strategy profile. We first argue that, in this case, the action profile $a^{11} \equiv (1,1,1,1,1,1)$ in the first period must also lead to a symmetric equilibrium payoff $w \in \mathcal{E}(\delta)$. To see this, note first that for the action profile a^{11} to be implementable in equilibrium it has to be enforceable on $\mathcal{E}(\delta)$. The convexity of $\mathcal{E}(\delta)$ implies it is equal to its convex hull so that the requirement is that we have to find three equilibrium payoffs $w_1, w_2, w_3 \in \mathcal{E}(\delta)$ such that (28) holds for each long-lived player. Since any equilibrium value $w \in \mathcal{E}(\delta)$ corresponds to a (unique) pair $(\mu_1, \mu_2) \in [0,1]$ such that $w = \hat{w}(\mu_1, \mu_2)$, finding three equilibrium values is equivalent to finding three pairs $\mu^0 = (\mu_1^0, \mu_2^0)$, $\mu^1 = (\mu_1^1, \mu_2^1)$, $\mu^2 = (\mu_1^2, \mu_2^2)$ in $M(\delta)$ such that (28) holds for each long-lived player with $w_j = \hat{w}(\mu^j)$.

With respect to player 1, (28) is

$$(1-\delta)u_L + \delta \sum_{j=0}^{2} \mathbb{P}\{s=j|a^{11}\}\hat{w}_1(\mu^j) \geq (1-\delta)(u_L + \lambda_L c) + \delta \sum_{j=0}^{2} \mathbb{P}\{s=j|0, a_{-1}^{11}\}\hat{w}_1(\mu^j). \quad (29)$$

With respect to player 2, (28) is

$$(1-\delta)u_L + \delta \sum_{j=0}^{2} \mathbb{P}\{s=j|a^{11}\}\hat{w}_2(\mu^j) \geq (1-\delta)(u_L + \lambda_L c) + \delta \sum_{j=0}^{2} \mathbb{P}\{s=j|0, a_{-2}^{11}\}\hat{w}_2(\mu^j).$$

With respect to player 3, (28) leads to three conditions of which, due to $\mathbb{P}\{s=j|01, a_{-3}^{11}\} = \mathbb{P}\{s=j|10, a_{-3}^{11}\}$, the latter two coincide:

$$(1-\delta)2u_G + \delta \sum_{j=0}^{2} \mathbb{P}\{s=j|a^{11}\}\hat{w}_3(\mu^j) \geq (1-\delta)2(u_G + \lambda_G c) + \delta \sum_{j=0}^{2} \mathbb{P}\{s=j|00, a_{-3}^{11}\}\hat{w}_3(\mu^j),$$

$$(1-\delta)2u_G + \delta \sum_{j=0}^{2} \mathbb{P}\{s=j|a^{11}\}\hat{w}_3(\mu^j) \geq (1-\delta)(2u_G + \lambda_G c) + \delta \sum_{j=0}^{2} \mathbb{P}\{s=j|10, a_{-3}\}\hat{w}_3(\mu^j),$$

$$(1-\delta)2u_G + \delta \sum_{j=0}^{2} \mathbb{P}\{s=j|a^{11}\}\hat{w}_3(\mu^j) \geq (1-\delta)(2u_G + \lambda_G c) + \delta \sum_{j=0}^{2} \mathbb{P}\{s=j|01, a_{-3}^{11}\}\hat{w}_3(\mu^j).$$

Combining and rewriting these conditions, we get local player 1's incentive constraint

$$\frac{1-\delta}{\delta}\frac{\lambda_L c}{(1-\alpha-\beta)u_L} \leq -\alpha\mu_1^0 + (2\alpha-1)\mu_1^1 + (1-\alpha)\mu_1^2; \quad (30)$$

local player 2's incentive constraint

$$\frac{1-\delta}{\delta}\frac{\lambda_L c}{(1-\alpha-\beta)u_L} \leq -\alpha\mu_2^0 + (2\alpha-1)\mu_2^1 + (1-\alpha)\mu_2^2; \quad (31)$$

the global player's incentive constraint not to shirk in both markets:

$$\frac{2(1-\delta)}{\delta}\frac{\lambda_G c}{(1-\alpha-\beta)u_G} \le -(1-\beta+\alpha)(\mu_1^0+\mu_2^0)+2(\alpha-\beta)(\mu_1^1+\mu_2^1)+(1-\alpha+\beta)(\mu_1^2+\mu_2^2); \quad (32)$$

and the global player's incentive constraint not to shirk in only one market:

$$\frac{1-\delta}{\delta}\frac{\lambda_G c}{(1-\alpha-\beta)u_G} \le -\alpha(\mu_1^0+\mu_2^0)+(2\alpha-1)(\mu_1^1+\mu_2^1)+(1-\alpha)(\mu_1^2+\mu_2^2). \quad (33)$$

The action profile a^{11} is implementable if a triple of pairs $\bar{\mu} = (\bar{\mu}^0, \bar{\mu}^1, \bar{\mu}^2)$ in $M(\delta)$ exist that together satisfy the constraints (30), (31), (32), (33). In this case, implementing a^{11} in the first period with continuation payoffs $\hat{w}^0(\mu^0)$, $\hat{w}^1(\mu^1)$, $\hat{w}^2(\mu^2)$ yields an aggregate payoff of

$$V(\mu^0,\mu^1,\mu^2) = (u_G+u_L)[2(1-\delta)+\delta\{\alpha^2(\mu_1^0+\mu_2^0)+2\alpha(1-\alpha)(\mu_1^1+\mu_2^1)+(1-\alpha)^2(\mu_1^2+\mu_2^2)\}] \quad (34)$$

Hence, the triple of pairs, $\bar{\mu} = (\bar{\mu}^0, \bar{\mu}^1, \bar{\mu}^2)$, that maximizes aggregate payoffs from an equilibrium strategy that implements a^{11} in the first period is a solution to the following program

$$\mathcal{P}(a^{11}) : \max_{\mu=(\mu^0,\mu^1,\mu^2)\in M(\delta)} V(\mu) \text{ s.t. } (30),(31),(32),(33).$$

If a solution to $\mathcal{P}(a^{11})$ exists, then there is one that is symmetric in the sense that $(\bar{\mu}_1^0, \bar{\mu}_1^1, \bar{\mu}_1^2) = (\bar{\mu}_2^0, \bar{\mu}_2^1, \bar{\mu}_2^2)$, because for any asymmetric solution $\bar{\mu}$ its symmetric average $\tilde{\mu}$ with $\tilde{\mu}_1^i = \tilde{\mu}_1^i = (\bar{\mu}_1^i + \bar{\mu}_2^i)/2$ has the same objective value V, lies in $M(\delta)$ (due to the convexity of $\mathcal{E}(\delta)$), and also satisfies all constraints (since the original $\bar{\mu}$ does so).

Using this observation, program $\mathcal{P}(a^{11})$ simplifies to finding three scalars (μ_a, μ_b, μ_c) with $(\mu_a, \mu_a), (\mu_b, \mu_b), (\mu_c, \mu_c) \in M(\delta)$ that maximize

$$W = (u_G+u_L)[2(1-\delta)+2\delta\{\alpha^2\mu_a+2\alpha(1-\alpha)\mu_b+(1-\alpha)^2\mu_c\}] \text{ s.t.} \quad (35)$$

$$\frac{1-\delta}{\delta}\frac{\lambda_L c}{(1-\alpha-\beta)u_L} \le -\alpha\mu_a+(2\alpha-1)\mu_b+(1-\alpha)\mu_c; \quad (36)$$

$$\frac{1-\delta}{\delta}\frac{\lambda_G c}{(1-\alpha-\beta)u_G} \le -(1-\beta+\alpha)\mu_a+2(\alpha-\beta)\mu_b+(1-\alpha+\beta)\mu_c; \quad (37)$$

$$\frac{1-\delta}{\delta}\frac{\lambda_G c}{2(1-\alpha-\beta)u_G} \le -\alpha\mu_a+(2\alpha-1)\mu_b+(1-\alpha)\mu_c. \quad (38)$$

Implementation of $a^{10} \equiv (1,0,1,0,1,0)$ in the first period We next study the implementability and optimality of the asymmetric action profile $a^{10} \equiv (1,0,1,0,1,0)$ in the first period. If the equilibrium that maximizes aggregate payoffs is such that it implements the asymmetric action $a^{10} \equiv (1,0,1,0,1,0)$ in the first period, then the equilibrium attains the value $V^* = V(\mu_1^* + \mu_2^*) = (u_G + u_L)(\mu_1^* + \mu_2^*)$. Moreover, it requires that the action a^{10} is enforceable in $\mathcal{E}(\delta)$. Hence, there must be three pairs $\mu^0 = (\mu_1^0, \mu_2^0)$, $\mu^1 = (\mu_1^1, \mu_2^1)$, $\mu^2 = (\mu_1^2, \mu_2^2)$ in $M(\delta)$ such that (28) holds for $a = a^{10}$ for each long-lived player with $w_j = \hat{w}(\mu^j)$. That is for player 1, we have

$$(1-\delta)u_L + \delta \sum_{j=0}^{2} \mathbb{P}\{s = j | a^{10}\} \hat{w}(\mu^j) \geq (1-\delta)(u_L + \lambda_L c) + \delta \sum_{j=0}^{2} \mathbb{P}\{s = j | 0, a_{-1}^{10}\} \hat{w}(\mu^j). \quad (39)$$

For player 2, we have

$$(1-\delta)0 + \delta \sum_{j=0}^{2} \mathbb{P}\{s = j | a^{10}\} \hat{w}(\mu^j) \geq -(1-\delta)\lambda_L c + \delta \sum_{j=0}^{2} \mathbb{P}\{s = j | 1, a_{-2}^{10}\} \hat{w}(\mu^j). \quad (40)$$

Since $\mathbb{P}\{s = j | a^{10}\} = \mathbb{P}\{s = j | 1, a_{-1}^{10}\}$ for all $j = 0, 1, 2$, constraint (40) is satisfied for any triple μ^0, μ^1, μ^2. For player 3, we have

$$(1-\delta)u_G + \delta \sum_{j=0}^{2} \mathbb{P}\{s = j | a^{10}\} \hat{w}(\mu^j) \geq (1-\delta)(u_G + \lambda_G c) + \delta \sum_{j=0}^{2} \mathbb{P}\{s = j | 00, a_{-3}^{10}\} \hat{w}(\mu^j), \quad (41)$$

$$(1-\delta)u_G + \delta \sum_{j=0}^{2} \mathbb{P}\{s = j | a^{10}\} \hat{w}(\mu^j) \geq (1-\delta)(u_G - \lambda_G c) + \delta \sum_{j=0}^{2} \mathbb{P}\{s = j | 11, a_{-3}^{10}\} \hat{w}(\mu^j), \quad (42)$$

$$(1-\delta)u_G + \delta \sum_{j=0}^{2} \mathbb{P}\{s = j | a^{10}\} \hat{w}(\mu^j) \geq (1-\delta)u_G + \delta \sum_{j=0}^{2} \mathbb{P}\{s = j | 01, a_{-3}^{10}\} \hat{w}(\mu^j). \quad (43)$$

The second inequality (42) holds for any triple (μ^0, μ^1, μ^2), since $\mathbb{P}\{s = j | a^{10}\} = \mathbb{P}\{s = j | 11, a_{-3}^{10}\}$ for each $j = 0, 1, 2$. Moreover, (41) implies (43), since $\mathbb{P}\{s = j | 00, a_{-3}^{10}\} = \mathbb{P}\{s = j | 01, a_{-3}^{10}\}$ for each $j = 0, 1, 2$. Hence, a^{10} is enforceable on $\mathcal{E}(\delta)$ if and only if we find a triple (μ^0, μ^1, μ^2) in $M(\delta)$ such that (39) and (41) hold.

Rewriting (39) yields the retailer's incentive constraint

$$\frac{1-\delta}{\delta} \frac{\lambda_L c}{(1-\alpha-\beta)u_L} \leq -(1-\beta)\mu_1^0 + (1-2\beta)\mu_1^1 + \beta\mu_1^2. \quad (44)$$

Rewriting (41) yields the global player's incentive constraint

$$\frac{1-\delta}{\delta} \frac{\lambda_G c}{(1-\alpha-\beta)u_G} \leq -(1-\beta)(\mu_1^0 + \mu_2^0) + (1-2\beta)(\mu_1^1 + \mu_2^1) + \beta(\mu_1^2 + \mu_2^2). \quad (45)$$

The aggregate payoff associated with the strategy profile a^{10} that is enforceable on $\mathcal{E}(\delta)$ by

(μ^0, μ^1, μ^2) is

$$W(\mu^0,\mu^1,\mu^2) = (u_G+u_L)[(1-\delta)+\delta\{(1-\beta)\alpha(\mu_1^0+\mu_2^0)+[(1-\beta)(1-\alpha)+\alpha\beta](\mu_1^1+\mu_2^1)+(1-\alpha)\beta(\mu_1^2+\mu_2^2)\}]$$

It follows that the equilibrium with the maximum aggregate payoffs that implements the asymmetric action a^{10} in the first period is a solution $(\hat{\mu}^0, \hat{\mu}^1, \hat{\mu}^2)$ to the following linear program:

$$\mathcal{P}^{10}: \max_{\mu^0,\mu^1,\mu^2 \in M(\delta)} \quad W(\mu^0,\mu^1,\mu^2) \text{ s.t. } (44), (45).$$

Denote the value of this program as $\hat{W}^{10} = W(\hat{\mu}^0, \hat{\mu}^1, \hat{\mu}^2)$.

If the equilibrium that maximizes aggregate payoffs is such that it implements the asymmetric action a^{10} in the first period then it holds $\hat{W}^{10} = V^*$.

Comparison of a^{11} and a^{10}. We first show that for δ small, the action profile a^{11} is optimal whenever it is implementable.

Lemma 6. *Suppose $\delta \leq 1/2$ and a^{11} is implementable. Then implementing a^{11} is optimal.*

Proof. If a^{10} is not implementable in that no combination $(\hat{\mu}^0, \hat{\mu}^1, \hat{\mu}^2)$ in $M(\delta)$ exists that satisfies (44) and (45), then the result follows trivially, since the only other implementable action profile $a^{00} = (0,0,0,0,0,0)$, which yields aggregate payoffs of 0, which is weakly less than any aggregates payoffs from a a triple of pairs $\bar{\mu} = (\bar{\mu}^0, \bar{\mu}^1, \bar{\mu}^2)$ in $M(\delta)$ that implements profile a^{11}.

So suppose a^{10} is implementable. We next demonstrate that for $\delta \leq 1/2$, the aggregate payoffs associated with any triple $(\hat{\mu}^0, \hat{\mu}^1, \hat{\mu}^2)$ in $M(\delta)$ that implements a^{10}, the aggregate payoffs are less than $2(u_G + u_L)(1-\delta)$, a lower bound on the payoffs of implementing a^{11} when it is implementable. To show this, consider the relaxed version of program $\mathcal{P}^{10}$ in which we disregard (44). Denoting this relaxed program as $\tilde{\mathcal{P}}^{10}$ and its value as $\tilde{W}^{10}$, it follows $\tilde{W}^{10} \geq \hat{W}^{10}$. The relaxed program has constraint (45) binding, since disregarding this constraint yields a solution with $\mu_1^0 + \mu_2^0 = \mu_1^1 + \mu_2^1 = \mu_1^2 + \mu_2^2$ which violates (45).

A binding (45) implies

$$(1-\beta)(\mu_1^0+\mu_2^0) = (1-2\beta)(\mu_1^1+\mu_2^1) + \beta(\mu_1^2+\mu_2^2) - \frac{1-\delta}{\delta}\frac{\lambda_G c}{(1-\alpha-\beta)u_G}. \tag{46}$$

Substituting out the expression $(\mu_1^0 + \mu_2^0)$, program $\tilde{\mathcal{P}}^{10}$ rewrites as

$$\max_{\mu^1,\mu^2 \in M(\delta)} (u_G+u_L)\left\{1-\delta+\delta\left[(1-\beta)(\mu_1^1+\mu_2^1)+\beta(\mu_1^2+\mu_2^2)-\alpha\frac{1-\delta}{\delta}\frac{\lambda_G c}{(1-\alpha-\beta)u_G}\right]\right\}.$$

The solution to this exhibits $\mu_1^1 + \mu_2^1 = \mu_1^1 + \mu_2^1 = \mu_1^* + \mu_2^*$ and, hence,[32]

$$\tilde{W}^{10} \leq (u_G + u_L) \left\{ 1 - \delta + \delta \left[(\mu_1^* + \mu_2^*) - \alpha \frac{1-\delta}{\delta} \frac{\lambda_G c}{(1-\alpha-\beta) u_G} \right] \right\}$$

Hence, $\tilde{W}^{10} \geq V^*$ implies

$$(u_G + u_L) \left\{ 1 - \delta + \delta \left[(\mu_1^* + \mu_2^*) - \alpha \frac{1-\delta}{\delta} \frac{\lambda_G c}{(1-\alpha-\beta) u_G} \right] \right\} \geq (u_G + u_L)(\mu_1^* + \mu_2^*)$$

so that

$$\mu_1^* + \mu_2^* \leq 1 - \alpha \frac{\lambda_G c}{(1-\alpha-\beta) u_G}$$

Hence, if the equilibrium that maximizes aggregate payoffs is such that it implements the asymmetric action $a^{10} \equiv (1, 0, 1, 0, 1, 0)$ in the first period, then we have $\mu_1^* + \mu_2^* \leq 1$. As a consequence, the maximum aggregate payoffs in $\mathcal{E}(\delta)$ is smaller than $(u_G + u_L)$, which for $\delta \leq 1/2$ is smaller than $2(u_G + u_L)(1 - \delta)$, a lower bound on the payoffs of implementing a^{11} when it is implementable. □

[32]The inequality is due to the fact that the right hand value is only attained if after substituting $\mu_1^1 + \mu_2^1 = \mu_1^1 + \mu_2^1 = \mu_1^* + \mu_2^*$ into (46) yields a $\mu^0 \in \mathcal{E}(\delta)$, otherwise the value $\tilde{W}^{10}$ is smaller than the RHS.

References

Abreu, D., D. Pearce, and E. Stacchetti (1990). "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." Econometrica 58: 1041-1063.

Andersson, F. (2002). "Pooling Reputations." International Journal of Industrial Organization 20(5): 715–730.

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

Bhattacharyya, S. and F. Lafontaine (1995). "Double-Sided Moral Hazard and the Nature of Share Contracts." RAND Journal of Economics 26:761–781.

Blair, R., and Lafontaine (2005). The Economics of Franchising. Cambridge University Press.

Biscarri, W., S. Zhao, and R. Brunner (2018) "A simple and fast method for computing the Poisson binomial distribution function." Computational Statistics & Data Analysis 122: 92–100.

Cabral, L. (2000). "Stretching Firm and Brand Reputation." RAND Journal of Economics 31(4): 658–673.

Cabral, L. (2009). "Umbrella Branding with Imperfect Observability and Moral Hazard." International Journal of Industrial Organization 27: 206–213.

Cai, H., and I. Obara (2009). "Firm Reputation and Horizontal Integration." RAND Journal of Economics 40(2): 340-363.

Calboli, I. (2007). "The Sunset of "Quality Control" in Modern Trademark Licensing." American University Law Review 57: 341–407.

Castriota S., and M. Delmastro (2012). "Seller Reputation: Individual, Collective, and Institutional Factors." Journal of Wine Economics 7: 49–69.

Choi, J. (1998). "Brand Extension and Information Leverage." Review of Economic Studies 65: 655–669.

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

Fleckinger, P. (2014). "Regulating Collective Reputation." Mimeo CERNA, MINES ParisTech.

Fudenberg, D., D. Levine, E. Maskin (1994). "The Folk Theorem with Imperfect Public Information." Econometrica 62 (5), 997-1039.

Hadfield, G. (1990). "Problematic Relations: Franchising and the Law of Incomplete Contracts." Stanford Law Review 42: 927–992.

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

Holmström, B. (1982). "Moral Hazard in Teams." The Bell Journal of Economics, 13: 324–340.

Klein, B., and K. Leffler, (1981). "The role of market forces in assuring contractual performance." Journal of Political Economy 89: 615–641.

Klein, B., and L. Saft. (1985). "The Law and Economics of Franchise Tying Contracts." The Journal of Law & Economics, 28: 345–361.

Kotler, P. (2003). Marketing Management. 11th Edition, Prentice-Hall, Upper Saddle River.

Levin, J. (2003). "Relational Incentive Contracts." American Economic Review 93 (3), 835-857.

Mailath, G., and L. Samuelson (2006). Repeated Games and Reputations: Long-Run Relationships. Oxford University Press.

Matsushima, H. (2001). "Multimarket Contact, Imperfect Monitoring, and Implicit Collusion." Journal of Economic Theory 98: 158–178.

Miklós-Thal, J. (2012). "Linking reputations through collective branding." Quantitative Marketing and Economics 10: 335–374.

Moorthy, S. (2012). "Can brand extension signal product quality?" Marketing science 31: 756-770.

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

Nosko, C., and S. Tadelis (2015). "The Limits of Reputation in Platform Markets: An Empirical Analysis and Field Experiment." Mimeo UC Berkeley.

Tirole, J. (1996). "A Theory of Collective Reputation (with Application to Corruption and Firm Quality)." Review of Economic Studies 63: 1–22.

Rukhin, A., C. Priebe, and D. Healy Jr. (2009). "On the monotone likelihood ratio property for the convolution of independent binomial random variables." Discrete Applied Mathematics 157: 2562–2564.

Wang, Y. (1993). "On the number of successes in independent trials." Statistica Sinica, 3: 295–312.

Wernerfelt, B. (1988). "Umbrella branding as a signal of new product quality: An example of signalling by posting a bond." The RAND Journal of Economics 30: 458–466.

Winfree, J., and J. McCluskey (2005). "Collective Reputation and Quality." American Journal of Agricultural Economics 87: 206–213.