Reputational Concerns with Altruistic Providers*

Pau Olivella† Luigi Siciliani‡

July 6, 2017

Abstract

We study a model of reputational concerns when doctors differ in their degree of altruism and they can signal their altruism by their (observable) quality. When reputational concerns are high, following the introduction or enhancement of public reporting, the less altruistic (bad) doctor mimics the more altruistic (good) doctor. Otherwise, either a separating or a semi-separating equilibrium arises: the bad doctor mimics the good doctor with probability less than one. Pay-for-performance incentive schemes are unlikely to induce crowding out, unless some dimensions of quality are unobservable. Under the pooling equilibrium a purchaser can implement the first-best quality by appropriately choosing a simple payment scheme with a fixed price per unit of quality provided. This is not the case under the separating equilibrium. Therefore, policies that enhance public reporting complement pay-for-performance schemes.

Keywords: reputation; altruism; doctors, name and shame policies, pay for performance.

JEL Classifications: I11; I18.

*The authors wish to thank the Editor, two anonymous referees, Frank Powell (who discussed our paper in the EHEW in Lausanne), and the participants in the EHEW, the Game Theory Workshop at the Institut d'Anàlisi Econòmica, the Winter Forum at the Barcelona GSE, and the Seminar at CHESEO (Oxford University) for their comments and suggestions. Any remaining errors are the authors' sole responsibility. Olivella acknowledges financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2015-0563), from the Spanish Ministry of Science and Education and FEDER, UE, through projects ECO2012-31962 and ECO2015-63679-P; and from the Spanish ONCE foundation.

†Department of Economics, Universitat Autònoma de Barcelona, Spain. E-mail: pau.olivella@uab.es; MOVE, CODE and Barcelona GSE.

‡Department of Economics and Related Studies, University of York, Heslington, York YO10 5DD, United Kingdom. E-mail: luigi.siciliani@york.ac.uk.

1 Introduction

A key policy issue in the health sector is how to incentivise providers (eg doctors, hospitals) to improve care. Providers have two sources of motivation: monetary and non-monetary ones. Monetary incentives include pay-for-performance incentive schemes: for example, hospitals are paid a price for each patient treated; family doctors can be financially rewarded if they have better quality indicators. Non-monetary incentives can be equally important and include two other incentive forces. First, providers may be altruistic and care about patients' well-being. Altruism motivates them to provide better quality and has long been recognised in the health economics literature (Ellis and McGuire, 1986; Chalkley and Malcomson, 1998). Second, providers care about what other people (their family, the community, their peers, other patients; henceforth society) think about them. Policymakers increasingly publish, and make available to patients and the general public, information on doctors' performance. Examples include the scheme "QualityCounts" in Wisconsin which compares adverse events in hospitals (Hibbard *et al.*, 2005); the Hospital Quality Alliance, which encourages public hospital reporting for a minimum of ten quality measures regarding three clinical conditions (Lindenauer *et al.*, 2007); and report cards for coronary bypass in Pennsylvania and New York State (Dranove *et al.*, 2003).[1] Such policies can potentially enhance reputational concerns by more widely advertising the performing doctors and the under-performing ones; they are sometimes (colloquially) known as *name and shame* schemes, where poorly performing doctors are subjected to shame in front of the community. Although reputational incentives have been recognised in the general economics literature (eg Benabou and Tirole, 2006 and 2011) we are not aware that they have been applied specifically to doctors and health care providers. This study fills this gap.

Can the simple fact of publishing information change doctors' behaviour? If so, which doctors change their behaviour and in which direction? Do patients and doctors gain from such policies? This study investigates the extent to which name and shame policies can enhance reputational concerns and induce some doctors to provide more quality to avoid a reputational damage. Health systems differ in the extent to which they compare and report quality in the public domain. They can vary from a small to a large set of indicators. They can report quality at organisation level (practice, hospital) or at individual doctor level, the latter exposing the doctors more directly. They can post indicators on a website, or more proactively disseminate the indicators by publishing them in newspapers. Variations in reporting generates variations in doctors' reputational concern. We investigate the effects of such variations in reputational concerns induced by different intensity of quality

[1] Analogous schemes have been implemented in other countries, sometimes in combination with pay-for-performance schemes, such as Brazil, Estonia, Korea, New Zealand and the United Kingdom (see Cashin *et al.*, 2014, p. 44-51).

reporting. In our model doctors differ in altruism and care about their own reputation: they enjoy being known by society as good doctors, and dislike being known as bad doctors. The analysis is distinct from the literature on financial incentives of healthcare providers since name and shame policies do not involve any direct payment, and instead require a signalling model to investigate how reputation is created in the first place.

We also investigate other policy relevant questions. First, we study whether a more extensive use of monetary incentives such as pay-for-performance schemes crowds out or crowds in the non-monetary incentives. Second, we investigate whether the benefits from publishing and disseminating information also arise within a multi-tasking framework when doctors provide different dimensions of quality, some of which are unobservable (eg diagnostic effort). Third, we investigate whether there is still scope for publishing and disseminating information on quality even when the purchaser (a health authority or a health insurer) can design a pay-for-performance scheme which pays a fixed price for each unit of quality provided. Since our focus is an applied one, we restrict the instruments which are available to the purchaser to *linear* performance contracts, since they are commonly used by policymakers.

Our model predicts that policies that publicise doctors' performance may be virtuous. By increasing reputational concerns, name and shame policies induce the bad doctor to mimic the higher quality provided by the good doctor. Whether the introduction of a pay-for-performance scheme crowds out or crowds in the non-monetary incentives is in principle indeterminate. Higher prices increase the good doctor's performance, and make more costly for the bad doctor to mimic the good doctor, which favours crowding out. But higher prices also increase overall revenues when performance is high, and make more attractive for the bad doctor to mimic the good doctor, which favours crowding in. We show that whether crowding in or crowding out arises ultimately depends on whether the good doctor provides *proportionally* lower or higher quality compared with the bad doctor in the absence of reputational payoffs. If the marginal benefit is decreasing, then under some regularity conditions on third-order derivatives of costs, the good doctor provides proportionally lower quality and *crowding in* arises.

Therefore, policies that introduce a pay-for-performance scheme do not seem to be in conflict with the introduction of report cards. However, this conclusion holds only if quality can be observed by patients and society. If some dimensions of quality cannot be observed (ie in the presence of multitasking), then name and shame policies can induce the bad doctor to *crowd out* non-observable dimensions of quality, and potentially reduce patients' benefit. Although crowding out is also found in the multitasking literature (Eggleston, 2005; Kaarboe and Siciliani, 2011), name and shame policies are not exempt from this issue: publishing in the public domain only a narrow set of quality indicators might make such policies undesirable. Moreover, from a modelling point of view, we show that multitasking interacts with reputational concerns by increasing the scope for the pooling equilibrium to arise. Name and shame policies make it easier for the bad doctor to mimic the good doctor by saving costs on the unobservable quality.

As for the optimal design of simple pay-for-performance schemes, we show that a linear contract which pays a fixed price per unit of quality is sufficient to achieve allocative efficiency for all doctors only if reputational concerns are high: the payer can design the incentive scheme aimed at the good doctor, and by pooling accomplishes efficiency of the bad doctor as well. This cannot arise for low reputational concerns unless more sophisticated non-linear contracts are available to the purchaser (eg the purchaser can offer a menu of contracts, which are not commonly observed in practice). Therefore, if the purchaser is constrained by the use of a linear performance contract, policies which publicize quality can make patients and purchasers better off even when the payment is optimally set. The result is relevant for policy and suggests that policies aimed at disseminating quality indicators have a role even in the presence of pay-for-performance schemes.

Our results are consistent with some empirical studies evaluating the effects of publicizing performance reports. Hibbard *et al.* (2005) compare the evolution of quality standards in obstetrics for (i) hospitals that had their reports made public; (ii) hospitals that received the report privately; and (iii) hospitals that did not receive any report. These authors find that "[a]mong the eight 'public report' hospitals with [...] low scores at baseline, only one had a worse-than-expected score two years later. In contrast, two-thirds of such hospitals in the 'private report' group and almost as many in the 'no report' group still had worse-than-expected scores two years later" (p. 1155). This suggests that hospitals with low quality responded to their reports being made public by improving performance. Similarly, Fichera *et al.* (2014) report in their survey that "[e]vidence from [the Hospital Quality Incentive Demonstration] and [the Advancing Quality] initiatives suggests that providers quickly converge to similar values on the process metrics and differences in performance must be measured at a very high level of precision to discriminate among providers." (p. 113) Wang et al. (2011) examine the impact of coronary bypass report cards. They find that poorly performing hospitals or surgeons responded with a reduction in volume, while highly rated hospitals and surgeons did not respond.

1.1 Related Literature

The empirical and theoretical literature on altruism and intrinsic motivation is extensive. Within the public and health economics literature the assumption of motivated agents is commonly shared.[2] Establishing that reputational concerns matter has also been in-

[2] Within the public economics literature, see Francois (2000), Besley and Ghatak (2005), Dixit (2005), Lakdawalla and Philipson, (2006), Delfgaauw and Dur (2008), Glazer (2004), Prendergast (2007), Makris

vestigated. Some empirical evidence has quantified the effects of publicizing performance indicators, either in isolation (Hibbard *et al.*, 2005) or when combined with other pay-for-performance schemes (Lindenauer *et al.*, 2007; see Roland and Dudley, 2015, for a review).

However, few studies formally include the possibility that reputational concerns come from society learning about doctors' altruism from observed actions. These studies can be classified into two groups. In the first group, this effect is either directly assumed in the doctor's payoff function (Siciliani, 2009) or comes about from the implicit assumption that individuals with different altruism choose different actions (Bénabou and Tirole, 2006). In the second group, reputational concerns are explicitly modelled as a formal signalling game. Bénabou and Tirole (2011) exploit a model which differs from ours in having a continuous type space and a dichotomous signal. Moreover, their focus is on analyzing and comparing sources of crowding out that are very different from ours. Much closer to our work are the studies by Jeitschko and Normann (2012), and Cartwright and Patel (2013). These authors consider signals with continuous support and dichotomous types. Jeitschko and Normann (2012), in an experimental game theory exercise, present a model that is close to ours but more restrictive. Cartwright and Patel (2013) present a model of fundraising and donations. However, they rule out the high-performance pooling equilibrium by means of the "Intuitive criterion" (Cho and Kreps, 1987), which leads them to focus on the so called "Riley Outcome" (Riley, 1979), which is a separating equilibrium by construction. We instead propose a simpler equilibrium selection procedure, based on more naive beliefs. In a nutshell, we first posit that doctors choose to perform as if no reputational concerns were present. Then we check whether they want to deviate from the posited choice in the presence of reputational concerns. We provide sufficient conditions under which our equilibrium outcome and beliefs pass the Intuitive Criterion test. We also provide sufficient conditions under which the Riley Outcome is not a perfect Bayesian equilibrium. Chen (2011) also investigates a signalling model where low-quality doctors select patients with low severity to mimic high-quality doctors. He shows that if doctors face the same distribution of patients' types, low-quality doctors have no incentive to select patients, while this arises when they face different distributions of patients' types. Rodriguez-Barraquer and Xu (2015) have agents seek promotion by choosing a difficult task. These authors also obtain a pooling equilibrium under some conditions.

Crowding out does arise in other settings. In Bénabou and Tirole (2006), financial rewards make it more difficult for society to infer types from observed actions (see also

(2009) and Makris and Siciliani (2013). Within the health economics literature the analytically-similar assumption of altruistic agents was introduced by Ellis and McGuire (1986), and then extended by Chalkley and Malcomson (1998), Eggleston (2005), Jack (2005), Siciliani (2009), Choné and Ma (2011), Brekke, Siciliani and Straume (2011, 2012), Kaarboe and Siciliani (2011), Siciliani, Straume and Cellini (2013) and Kolstad (2013).

Ariely *et al.*, 2009). In psychology, in particular in Self Determination Theory, Ryan and Deci (2000) propose that tangible rewards suppress the direct pleasure that one's actions produce. In cases where the true objectives of the payer (or some aspect of the environment) are unknown to agents, the mere fact that the principal introduces extrinsic motivation may in itself (partly) reveal such information (see, for instance, Fehr and List, 2004; Falk and Kosfeld, 2006; Funk, 2007; Bénabou and Tirole, 2011). This effect is absent in our analysis since doctors do not care about the payer's type and are fully aware of the impact their actions have on patients' well-being.

In Section 2 we present the model and characterize the equilibria. In Section 3 we investigate if pecuniary incentives generate crowding out or crowding in. In Section 4 we extend the model to a multitasking environment. In Section 5 we characterize the optimal (linear) remuneration contract. In Section 6 we provide conditions for our equilibrium to pass the Intuitive Criterion test and for the Riley outcome not being an equilibrium. Section 7 discusses some extensions, and Section 8 concludes. Technical derivations are relegated to the Appendix.

2 The model

The players are a doctor, a third-party payer (a public or private insurer), and society (patients, their family and friends, doctor's peers). Define q as the quality of care received by the patients. Doctors are altruistic and differ in the degree of altruism, which can take two values $\theta \in \left\{\underline{\theta};\overline{\theta}\right\}$ with $\underline{\theta} < \overline{\theta}$, and is private information. We refer to the more altruistic provider as the *good doctor*, and to the less altruistic provider as the *bad doctor*. The prior probability that the doctor is good is common knowledge and equal to $\lambda \geq 0$. Both types have the same costs of delivering quality, $C(q)$, with $C_q > 0$ and $C_{qq} \geq 0$. These costs are the sum of the monetary and non-monetary ones (eg diagnostic effort and opportunity costs of time spent with the patient).

Although we interpret q as quality, it can be interpreted more broadly as intensity of care. Under the latter interpretation we allow the marginal benefit of q to become negative for high intensity of care (as in Ellis and McGuire, 1986). Formally, patients derive benefits $W(q)$, with $W_q(q) \geq 0$ if $q \leq \widehat{q}$, $W_q(q) \leq 0$ if $q \geq \widehat{q}$, and $W_{qq} \leq 0$. The marginal benefit is negative if unnecessary tests, X-rays, or drugs with side effects and no health gains are prescribed. For brevity we refer to q as quality in the rest of the paper.

Patients observe quality and use that observation to update their beliefs on the doctor's type, ie to decide whether the doctor is good. We denote these (posterior) beliefs as λ^S, which is the probability that the doctor is good after having observed quality. We denote the expected type of a doctor by $\theta^S = \lambda^S\overline{\theta} + \left(1-\lambda^S\right)\underline{\theta}$. If there is no updating, then the expected type is the expected type in doctors' population ($\lambda^S = \lambda$). If updating is such that the doctor is good (bad), then the expected type is $\bar{\theta}$ ($\underline{\theta}$) and $\lambda^S = 1$ ($\lambda^S = 0$).

Doctor's preferences are represented by a linear and additively separable utility function over money, altruism and reputation. The revenues are given by $T + pq$, where T is a fixed budget (or lump-sum payment) and p is a bonus for additional quality (eg as part of a pay-for-performance scheme).[3] His profits are $\pi(q) = T + pq - C(q)$. We also assume that doctors have limited liability and that the purchaser cannot use a contract (or a menu of contracts) that specifies a transfer conditioned on a given quality (see also Section 5 on optimal contracting).

Similarly to Ellis and McGuire (1986) and Chalkley and Malcomson (1998), altruism is expressed as a fraction of patients' benefits, $\theta W(q)$, and the doctor cares about patients' wellbeing (Andreoni, 1989). θ can alternatively be interpreted as intrinsic motivation (Dixit, 2005; Besley and Ghatak, 2005). The sum of the first two components gives the *non-reputational payoff* defined as

$$V(q) = \pi(q) + \theta W(q). \tag{1}$$

The reputational concerns convey that the doctor cares about society's *impression* of his own altruism, and come from the composition of two elements. The first element $\alpha \equiv \alpha_0 (\alpha_1 + \alpha_2)$ measures how intensely society, or the doctor himself, takes into account this impression. The parameter α_0 is determined by society's preferences (how much society cares about altruism) or the doctor's own preferences (how much he cares about what others think of him). α_1 measures the number of people who directly learn, without any policy intervention, about the quality of his care through family, friends, word of mouth, and social networks.

α_2 is a key policy parameter and captures the extent to which policymakers can change and amplify reputational concerns by publishing and disseminating quality reports. Low levels of α_2 are associated with health systems with limited or no quality reports; reporting is at higher organisation level; and there are limited efforts to disseminate information. In contrast, high levels of α_2 are associated with health systems which publish an extensive range of quality indicators; the reporting is at the individual doctor level, which exposes

[3] Examples include the Medicare Programme in the United States, which financially rewards hospitals that do well according to measurable quality indicators, such as rates of cervical cancer screening and haemoglobin testing for diabetic patients (Rosenthal *et al.*, 2005). In the United Kingdom, general practitioners performing well on certain quality indicators, such as the measurement of blood pressure and cholesterol in patients with ischemic heart disease, can receive substantial financial rewards (up to 20% of revenues, Doran *et al.*, 2006). Hospitals receive Best Practice Tariffs for a selection of conditions, such as hip fracture and stroke. An additional payment is provided, on top of a basic DRG tariff, conditional on performance related to a process measure of quality (e.g. rapid brain imaging or being treated in a stroke unit). Rosenthal *et al.* (2004) provide 36 other examples of pay-for-performance programs in the United States. Similar initiatives are under discussion in Australia, Canada, New Zealand, the Netherlands and Spain (Gravelle, Sutton and Ma, 2010).

them more directly; and the quality reports are widely disseminated through websites, newspapers, and leaflets. The aim of this section is to investigate how changes in the policy design of public reporting affects doctor's behaviour through changes in reputational concerns.

The second element of the reputational concerns is how society's impression is determined. This is given by the difference between the conditional expectation of altruism given an observed quality, $\theta^S(q)$, and the expected altruism, $E(\theta) = \lambda\bar{\theta} + (1-\lambda)\underline{\theta}$. It is consistent with the idea that if no new information is revealed, the doctor's reputation remains the same. If all doctors provide the same quality, quality is not informative and there is no reputational gain or loss. If observing some (low) quality generates posterior beliefs that the doctor is bad, then the reputational payoff is negative. To sum up, doctors' preferences are

$$V(q) + \alpha(\theta^s(q) - E(\theta)). \tag{2}$$

We define the *non-reputational optimum quality* $q^*(\theta)$ as the optimal quality of doctor θ facing no reputational concerns (when $\alpha = 0$), which maximises $V(q \mid \theta)$ and satisfies $p + \theta W_q = C_q$, so that the marginal benefit from quality due to monetary and altruistic concerns is equal to the marginal cost.

Our equilibrium concept is the Perfect Bayesian Equilibrium (PBE, see Appendix 1 for a formal definition). As it is usually the case, the fact that the PBE notion does not restrict beliefs for actions out of equilibrium leads to a plethora of PBE. We therefore restrict beliefs to satisfy the following two properties: (1) [Monotonicity] Beliefs on altruism are not decreasing in quality; (2) [Pessimism] Beliefs for out-of-equilibrium actions are the most pessimistic. Monotonicity and Pessimism imply that society does not raise its beliefs that the doctor is good when observing quality levels that are higher than the quality chosen by the good doctor in equilibrium.[4] Even imposing monotonicity and pessimism, there is still a large multiplicity of PBEs. The literature has often used the Intuitive Criterion (Cho and Kreps, 1987) to restrict out of equilibrium beliefs. Rather than applying this criterion directly, we focus on equilibria where the good doctor chooses the quality which maximises his non-reputational payoff, and the bad doctor may mimic such quality. In Section 6 we provide conditions under which this equilibrium passes the Intuitive Criterion test. We now characterise the parameter values which sustain a separating, a pooling and a semi-separating equilibrium, respectively.

Separating equilibrium. The good doctor provides the higher quality, $q^E(\bar{\theta}) = q^*(\bar{\theta})$ and the bad doctor provides the lower quality $q^E(\underline{\theta}) = q^*(\underline{\theta})$. Beliefs are such that observing a high (low) quality signals with certainty high (low) altruism. The good doctor never has an incentive to mimic the bad doctor, while the bad doctor has no incentive to

[4] Monotonicity in beliefs does not imply that the benefit function is monotonic and increasing in quality.

mimic the good doctor if the reputational gains are low (Appendix 1):

$$0 \leq \alpha \leq \tau = \frac{V(q^*(\underline{\theta}) \mid \underline{\theta}) - V(q^*\left(\overline{\theta}\right) \mid \underline{\theta}))}{\overline{\theta} - \underline{\theta}}. \tag{3}$$

The parameter τ has an intuitive interpretation: it conveys the cost of the bad doctor of disguising as a good doctor. In terms of payoffs, the good doctor enjoys an increase in reputation and has the payoff $V(q^*(\overline{\theta}) \mid \overline{\theta}) + \alpha(\bar{\theta} - E(\theta))$, while the bad doctor suffers a loss of reputation and has the payoff $V(q^*(\underline{\theta}) \mid \underline{\theta}) - \alpha(E(\theta) - \underline{\theta})$.

Pooling equilibrium. Suppose that both doctors choose the high quality: $q^E(\theta) = q^*(\overline{\theta})$ for all θ. Since doctors provide the same quality, patients (and society) cannot distinguish between good and bad doctors. There is therefore no updating in beliefs about altruism after observing the high quality, and the expected altruism is the average one. The good doctor never has an incentive to mimic the bad doctor. The bad doctor has an incentive to mimic the good doctor if the reputational gains are high (Appendix 1): $\alpha \geq \tau/\lambda$. No type enjoys a reputation gain or loss, and equilibrium payoffs are $V(q^*(\overline{\theta}) \mid \overline{\theta})$ for the good doctor and $V(q^*(\overline{\theta}) \mid \underline{\theta})$ for the bad doctor. This is unsurprising since observing the high quality is not informative. The bad doctor suffers a mimicking cost, which he is willing to incur to avoid a reputation loss.

Semi-separating equilibrium. Finally, there is an empty intersection for intermediate levels of reputational concerns, where neither the separating equilibrium nor the pooling equilibrium exists. This empty intersection is due to the impossibility under the pooling equilibrium for the bad doctor to capture the highest reputation gain. He must content himself with avoiding a reputation loss. A semi-separating equilibrium arises for $\tau < \alpha < \tau/\lambda$. The good doctor chooses the high quality $q^E\left(\overline{\theta}\right) = q^*\left(\overline{\theta}\right)$. The bad doctor chooses the low quality $q^*(\underline{\theta})$ with probability $r^E = 1 - [(\lambda/(1-\lambda))((\alpha-\tau)/\tau)] < 1$ and the high quality $q^*\left(\overline{\theta}\right)$ with probability $(1 - r^E)$.[5] (See Appendix 1).

We summarise all results obtained so far with the following proposition.

Proposition 1 *If reputational concerns are small (for $0 \leq \alpha \leq \tau$), a separating equilibrium arises where the good doctor provides a higher quality than the bad doctor, $q^E\left(\overline{\theta}\right) = q^*\left(\overline{\theta}\right) > q^E(\underline{\theta}) = q^*(\underline{\theta})$. The good doctor enjoys a reputation gain while the bad doctor suffers a reputation loss. If reputational concerns are high (for $\alpha \geq \tau/\lambda$), a pooling equi-*

[5] The equilibrium payoffs for the good and bad doctor are, respectively, $V(q^*\left(\bar{\theta}\right) \mid \underline{\theta}) + (\tau - \alpha\lambda)\left(\bar{\theta} - \underline{\theta}\right)$ and $V(q^*(\overline{\theta}) \mid \overline{\theta}) + (\tau - \alpha\lambda)\left(\bar{\theta} - \underline{\theta}\right)$. Hence the observation of low quality induces the sure belief that the provider is a bad doctor and the consequent reputation loss. In contrast, the observation of high quality is not fully informative. It could either come from a good doctor or a bad doctor mimicking the good doctor. Upon the observation of high quality, the posterior probability that the doctor is good is larger than the unconditional one (λ), and both doctors enjoy a reputation gain, which is small and equal to $(\tau - \alpha\lambda)\left(\bar{\theta} - \underline{\theta}\right)$, and decreasing in reputational concerns α. The latter is not surprising since the bad doctor tends to mimic with a higher probability when reputational concerns are larger.

librium arises where both doctors choose a high quality, $q^E\left(\overline{\theta}\right) = q^E\left(\underline{\theta}\right) = q^*\left(\overline{\theta}\right)$, *and neither type gains or loses any reputation. If reputational concerns are intermediate (for* $\tau < \alpha < \tau/\lambda$*), a semi-separating equilibrium arises where the good doctor provides the high quality and the bad doctor randomizes between high and low quality. In the latter equilibrium, the expected quality of the bad doctor increases with the intensity of reputational concerns. Payoffs decrease with the intensity of reputational concerns.*

The proposition generates the following key policy insight. Consider a policy which amplifies reputational concerns from low to high. For example, before the policy i) there is limited availability or no quality reporting; ii) the reporting, if it occurs, is at a higher organisation level so that doctors do not feel directly exposed to this information; and iii) the reports are not widely disseminated. After the policy, i) there is a plethora of quality reports; ii) the quality reports are provided at doctor level; and iii) quality reports are widely disseminated through websites, newspapers, and leaflets. Our analysis suggests that name and shame policies will have the intended effect of raising quality.

We conclude by further characterising the three equilibria. The semi-separating equilibrium smoothly connects the separating with the pooling equilibrium (Appendix 1), which is illustrated in Figures 1 (qualities) and 2 (payoffs). If reputational concerns are intermediate, the bad doctor increases quality in expected terms when reputation concerns are higher to avoid the increasingly important consequences of a bad reputation. In terms of payoffs, both the bad and the good doctors are worse off under a pooling equilibrium than either a separating or a semi-separating equilibrium.[6] Figure 1 also shows that when the proportion of good doctors λ is high (which is consistent with some experimental evidence on doctors' altruism; Godager and Wiesen, 2013), then name and shame policies are more likely to induce the bad doctor to increase quality driven by the higher expected altruism (in doctor population) and reputational gains.

[Figures 1 and 2 here]

[6] When reputational concerns are low, the payoff of the good (bad) doctor increases (decreases) with the intensity of reputational concerns. This is intuitive. Higher intensity amplifies the positive (negative) reputation payoff of being known as the good (bad) doctor. In contrast, when the intensity of reputation concerns are intermediate, both doctors' payoffs reduce with the intensity of reputation concerns. As the bad doctor increases quality (in expected terms), patients and society increasingly cannot distinguish between the two types and therefore the reputational payoff vanishes. When reputational concerns are high, patients and society cannot distinguish at all between good and bad doctors. The good doctor ends up with the same payoff obtained in the absence of reputational concerns. The bad doctor obtains the lowest payoff, being induced to exert the higher quality and being able only to avoid a bad reputation, instead of gaining a good one.

3 Crowding in or crowding out?

In this section, we investigate whether an increase in pecuniary incentives crowds out the effect of the reputational concerns. In our model this occurs if a change in price, due for example to the introduction of a 'pay-for-performance' scheme, affects the range of reputational concerns over which the pooling equilibrium arises. In turn, this depends on τ, ie the cost to the bad doctor of mimicking the good doctor.

Does an increase in price generate crowding out or crowding in? We obtain (Appendix 2):

$$\frac{\partial \tau}{\partial p} = \underbrace{-\frac{q^*(\overline{\theta}) - q^*(\underline{\theta})}{\overline{\theta} - \underline{\theta}}}_{\text{revenue effect}} + \underbrace{W_q(q^*(\overline{\theta}))\frac{\partial q^*(\overline{\theta})}{\partial p}}_{\text{quality effect}}. \tag{4}$$

The effect is indeterminate. An increase in price affects the revenues of the bad doctor more when he provides high quality than when he provides low quality (first term in (4)). This effect, which we call the *revenue* effect, is negative and tends to induce crowding in. However, a higher price also increases the qualities that maximise the non-reputational payoffs of the good doctor and the bad doctor, $q^*(\underline{\theta})$ and $q^*(\overline{\theta})$. By the envelope theorem we know that an increase in the lower quality $q^*(\underline{\theta})$ has no effect on the non-reputational payoff of the bad doctor. An increase in the higher quality $q^*(\overline{\theta})$ reduces it because it brings the bad doctor even further away from his non-reputational optimum quality (second term in (4)). This effect, which we call the *quality* effect, is positive and tends to induce crowding out. Which of the two effects dominates ultimately depends on whether the good doctor *proportionally* provides higher or lower quality than the bad doctor, which in turn depends on the sign of $\partial^2 q^*(\theta)/\partial^2\theta$ (Appendix 2).

Proposition 2 *Crowding in (out) arises if the good doctor proportionally provides lower (higher) quality than the bad doctor, ie if $\partial^2 q^*(\theta)/\partial^2\theta < 0$ (> 0).*

If the marginal benefit is decreasing, then under some regularity conditions on costs, the good doctor provides proportionally lower quality and therefore crowding in arises. A sufficient condition is that $W_{qq} < 0$, $W_{qqq} \leq 0$ and $C_{qqq} \geq 0$ (Appendix 2). This condition is satisfied if the cost of quality is, for instance, quadratic or linear. For crowding out to arise, the marginal benefit from quality has to be constant or mildly decreasing and the marginal cost has to be concave in quality.

The key policy insight is that under a wide range of scenarios crowding out is unlikely to arise. Therefore, policies that introduce a pay-for-performance scheme do not seem to be in conflict with the introduction of report cards. However, this conclusion is valid when report cards and pay-for-performance schemes capture all dimensions of treatment' quality. This assumption is relaxed in the next section.

4 Multitasking

In this section we extend the model to allow for multiple dimensions of quality, which leads to the *multitasking* problem: incentivizing one dimension of quality may trigger a reduction in the unincentivised dimension of quality (Holmstrom and Milgrom, 1991; Eggleston, 2005; Kaarboe and Siciliani, 2011). We assume that quality 1, q_1, can be observed by society, and quality 2, q_2, cannot be observed either because the patient lacks the expertise to evaluate it or because quality 2 is not reported. Both doctors have the same costs, $C(q_1, q_2)$, with $C_{q_1} > 0$, $C_{q_2} > 0$, $C_{q_1q_1} > 0$ and $C_{q_2q_2} > 0$. We assume qualities are *substitutes,* $C_{q_1q_2} > 0$ with $C_{q_1q_1}C_{q_2q_2} > C^2_{q_1q_2}$, and an increase quality 1 increases the marginal cost of quality 2. Patients benefit from both quality dimensions: $W(q_1, q_2)$, with $W_{q_1} > 0$, $W_{q_2} > 0$ and $W_{q_1q_2} = 0$. Doctor's profits are $\pi(q_1, q_2) = T + pq_1 - C(q_1, q_2)$, where p is a unit price for observable quality 1. The non-reputational payoff is $V = \pi + \theta W$. Reputational concerns arise only as a result of changes in observable quality 1 generating a payoff equal to $\alpha(\theta^s(q_1) - E(\theta))$. The *non-reputational* optimal qualities satisfy $p + \theta W_{q_1}(q_1^*) = C_{q_1}(q_1^*, q_2^*)$, $\theta W_{q_2}(q_2^*) = C_{q_2}(q_1^*, q_2^*)$: the marginal benefit of quality due to monetary and altruistic concerns equals the marginal cost. An increase in price incentivises the observable quality but disincentivises unobservable quality, $\partial q_1^*/\partial p > 0, \partial q_2^*/\partial p < 0$. Qualities increase in altruism, $\partial q_1^*/\partial\theta > 0, \partial q_2^*/\partial\theta > 0$, if the degree of cost substitution is not too high (all proofs in Appendix 3).

Pooling equilibrium. The good doctor chooses qualities which maximise the non-reputational payoff: $q_1^E(\bar{\theta}) = q_1^*(\bar{\theta})$, $q_2^E(\bar{\theta}) = q_2^*(\bar{\theta})$, and the bad doctor chooses the same observable quality 1 as the good doctor: $q_1^E(\underline{\theta}) = q_1^*(\bar{\theta})$. But since qualities are substitutes, the bad doctor compensates by providing a low unobservable quality, $q_2^M(\underline{\theta})$ satisfying $\underline{\theta} W_{q_2}(q_2^M) = C_{q_2}(q_1^*(\bar{\theta}), q_2^M)$, which is lower than under a *separating* equilibrium, $q_2^E(\underline{\theta}) = q_2^M(\underline{\theta}) < q_2^*(\underline{\theta})$. The bad doctor has an incentive to mimic the good doctor when reputation concerns are high:

$$\alpha \geq \frac{\tau_1}{\lambda} \overset{def}{=} \frac{1}{\lambda} \frac{V(q_1^*(\underline{\theta}), q_2^*(\underline{\theta}) \mid \underline{\theta}) - V(q_1^*(\bar{\theta}), q_2^M(\underline{\theta}) \mid \underline{\theta})}{\bar{\theta} - \underline{\theta}}. \tag{5}$$

Separating equilibrium. Each doctor chooses qualities $q_1^E(\theta) = q_1^*(\theta)$ and $q_2^E(\theta) = q_2^*(\theta)$ that maximise non-reputational payoffs. The good doctor never has an incentive to mimic the bad doctor. The bad doctor has no incentive to mimic the good doctor if reputational gains are low: $0 \leq \alpha \leq \tau_1$.

We summarise with the following proposition.[7]

[7] In line with Section 2, a *semi-separating* equilibrium arises for intermediate reputational concerns, i.e. $\tau_1 \leq \alpha \leq \tau_1/\lambda$, where the good doctor always provides high qualities in both dimensions, and the bad doctor randomises between the qualities provided under the separating and the pooling equilibrium. Higher reputational concerns increase the probability that the bad doctor provides the high observable

Proposition 3 *If reputational concerns are small (for $0 \leq \alpha \leq \tau_1$), a separating equilibrium arises where both types of doctor chose qualities which maximise the non-reputational payoff. If reputational concerns are high (for $\alpha \geq \tau_1/\lambda$), a pooling equilibrium arises where the good doctor chooses qualities which maximise the non-reputational payoff. Driven by reputational concerns, the bad doctor chooses the same observable quality chosen by the good doctor but compensates by providing a low level of the unobservable quality (which is lower than the unobservable quality maximising the non-reputational payoff).*

We conclude by highlighting differences with the existing literature on multitasking and key policy implications. The presence of multitasking changes the nature of the pooling equilibrium, which is at the core of this study. It increases the scope for the pooling equilibrium to arise and for policies which publish quality indicators to induce a change in doctors' behaviour. The presence of multitasking makes it easier for the bad doctor to mimic the good doctor, since the additional cost of increasing the observable quality can be offset by reducing the unobservable quality. In terms of policy, it is precisely this offsetting behaviour which makes publishing information less desirable. There may be scenarios in which the patients' benefit from unobservable quality is more important than the benefit from observable quality, so that publishing information may be harmful to patients. Formally, the comparison depends on $W(q_1^*\left(\overline{\theta}\right), q_2^M\left(\underline{\theta}\right)) \gtrless W(q_1^*\left(\underline{\theta}\right), q_2^*\left(\underline{\theta}\right))$. Although this result is in line with the multitasking literature, from a policy perspective it is important to highlight that name and shame policies are not exempt from this issue and can be undesirable as a result. Moreover, the introduction of a pay-for-performance scheme may exacerbate the multitasking problem when reputational concerns are high. An increase in price will increase observable quality and it will further reduce unobservable quality. Although this also arises for low reputation concerns, unobservable quality is always lower under a pooling equilibrium than under a separating equilibrium. Therefore, under the assumption of decreasing marginal benefit of the unobservable quality, it is likely that the patient will suffer more from a reduction in unobservable quality when reputational concerns are high.

Our results are consistent with some empirical evidence. Lu (2012) finds that after the introduction of the Nursing Home Quality Initiative in the US "scores of quality measures improve for the reported dimensions but deteriorate for the unreported ones". Our model also predicts that name and shame policies will reduce the variance in observable quality and increase the variance in unobservable quality. This is also in line with the empirical evidence. Werner et al. (2009) analyse the effects of public reporting in nursing facilities and find that, after the publication of some measures of quality, "[f]or unreported measures that worsened on average, the decrement in quality was generally larger among low-scoring

quality and the low unobservable quality.

facilities than high-scoring facilities" (p.390), whereas "all the three reported measures of quality improved after the reporting policy was implemented"(p.388).

5 Optimal contracting

In this section we relax the assumption that the doctor's contract is exogenous, and allow the purchaser of health services (eg a public or private insurer) to design the optimal contract by endogenising the choice of the pay-for-performance price p and of the fixed transfer T. We derive the optimal contract under the assumption that the purchaser is constrained to the use of the linear contract $T + pq$ which is independent of the doctor's type: both the good and the bad doctor receive the same fixed budget and the same price per unit of quality provided. Doctors can, however, differ in the quality provided. Such contracts are common in the health sector (Rosenthal et al., 2005; and Section 2). We solve by backward induction. We start by deriving the optimal price under a pooling equilibrium, and then verify ex-post the range of reputational concerns over which pooling arises when evaluated at such optimal price. We proceed similarly for the separating equilibrium.

Pooling equilibrium. With high reputational concerns both doctors provide the same high quality $q^*\left(\overline{\theta}, p\right)$ for a given price. We assume that the purchaser maximises the difference between patient's benefit and the transfer to the provider:

$$\underset{T,p}{Max}\ W\left(q^*\left(\overline{\theta}, p\right)\right) - T - pq^*\left(\overline{\theta}, p\right), \tag{6}$$

subject to two types of constraint: the participation constraints ensure that each doctor is willing to provide services no matter what his type, and the limited-liability constraints ensure that each doctor does not make a negative profit. The participation constraint is always satisfied when the limited-liability constraint is satisfied, since the utility is the sum of the profits and the altruistic component. Therefore, the limited-liability constraint is the only binding one (Appendix 4), which substituted into (6) gives the optimal contract:

$$p^P = \left(1 - \overline{\theta}\right) W_q\left(q^o\right), \quad T^P = C(q^o) - p^P q^o, \tag{7}$$

where quality q^o is such that it equates the marginal benefit with the marginal cost of quality, $W_q\left(q^o\right) = C_q\left(q^o\right)$. A higher level of altruism for the good doctor implies that the purchaser can set a lower price to induce the desired level of quality. The fixed transfer covers the differences between the costs and the revenues from the pay-for-performance scheme to ensure that the doctor breaks even. The key insight is that under a pooling equilibrium, the purchaser can obtain allocative efficiency by designing a contract that is aimed at the good doctor. Since the bad doctor mimics the good doctor, the same contract also induces allocative efficiency for the bad doctor. This pooling equilibrium arises for $\alpha \geq \tau(p^P)/\lambda$.

Separating equilibrium. With low reputational concerns, the purchaser problem is:

$$\underset{T,p}{Max}\ \lambda\left[W\left(q^*\left(\overline{\theta},p\right)\right)-T-pq^*\left(\overline{\theta},p\right)\right]+\left(1-\lambda\right)\left[W\left(q^*\left(\underline{\theta},p\right)\right)-T-pq^*\left(\underline{\theta},p\right)\right] \quad (8)$$

subject to the participation constraints and the limited-liability constraints. Since profit decreases in altruism, it is the liability-constraint of the good doctor that is binding, so that $T = C(q^*\left(\overline{\theta},p\right)) - pq^*\left(\overline{\theta},p\right)$. The participation constraint of the good and the bad doctor is never binding (Appendix 4). Substituting for T in (8) and maximising with respect to price, we obtain

$$\begin{aligned}&\lambda W_q\left(q^*\left(\overline{\theta},p\right)\right)\frac{\partial q^*(\overline{\theta},p)}{\partial p}+\left(1-\lambda\right)W_q\left(q^*\left(\underline{\theta},p\right)\right)\frac{\partial q^*\left(\underline{\theta},p\right)}{\partial p} \qquad (9)\\ =\ & C_q(q^*\left(\overline{\theta},p\right))\frac{\partial q^*(\overline{\theta},p)}{\partial p}-\left(1-\lambda\right)\left[q^*\left(\overline{\theta},p\right)-q^*\left(\underline{\theta},p\right)+p\left(\frac{\partial q^*(\overline{\theta},p)}{\partial p}-\frac{\partial q^*\left(\underline{\theta},p\right)}{\partial p}\right)\right].\end{aligned}$$

The optimal price is set such that the average marginal benefit (weighted by the response of quality to price) is equal to the marginal cost. The marginal cost has two components: since the limited-liability constraint is binding for the good doctor, the first term refers to the marginal cost of the good doctor (weighted by his responsiveness of quality to price). The second term accounts for rent extraction distortions: since the bad doctor makes a positive profit, it is optimal to distort prices to reduce such rents. The rent extraction term pushes the price upwards if doctors do not differ significantly in their quality responsiveness to price. The optimal contract is given by the pair $\left\{p^S, T^S\right\}$ where p^S denotes the optimal price under the separating equilibrium in (9) and $T^S = C(q^*\left(\overline{\theta},p^S\right)) - p^S q^*\left(\overline{\theta},p^S\right)$ is the optimal lump-sum transfer to ensure the doctors break even. The separating equilibrium arises for $\alpha \leq \tau(p^S)$.

The key policy insight is that when the purchaser is constrained to the use of the linear contract, $T+pq$, the purchaser is better off under a pooling equilibrium than under a separating equilibrium. This is because reputational concerns reduce (eliminate) differences in qualities between different doctor types. In turn, the purchaser can implement allocative efficiency for both types by setting a price which is targeted at the good doctor. The bad doctor simply mimics the good doctor. Neither type makes a profit. In contrast, the purchaser is constrained under a separating equilibrium. Since different types provide different qualities for a given price, the purchaser aims at inducing allocative efficiency for the average type. Moreover, it distorts price to reduce informational rents for the bad doctor. If differences in responsiveness of quality to price between the good and the bad doctor are small, then the optimal price is higher under a separating equilibrium than under a pooling equilibrium. This arises for two main reasons: first, since the average type is less motivated than the good type, the purchaser needs to incentivise more doctors through a higher-powered incentive scheme; second, a higher price helps to reduce the informational rent of the bad doctor. The purchaser is also better off under a pooling equilibrium than under a semi-separating one, $\tau(p^S) \leq \alpha \leq \tau(p^P)/\lambda$ (Appendix 4), and the optimal price under a semi-separating equilibrium is between the price under pooling and under separating.

We summarise in the following proposition.

Proposition 4 *If the purchaser is constrained to the use of the linear contract $\{T, p\}$ which is independent of doctor's type, high reputational concerns induce a pooling equilibrium where both doctors provide the same quality and make the purchaser better off.*

Name and shame policies and pay-for-performance schemes tend to be substitutes. Policymakers can induce doctors to provide higher quality either by publishing and disseminating performance indicators or by tying the performance indicators to financial incentives. This section has shown that if the purchaser is constrained to a linear contract, the purchaser may do better by widely disseminating information than by strengthening pay-for-performance schemes. However, the result critically relies on the assumption that policy makers are constrained to linear contracts. It can be shown that if the purchaser can implement more flexible non-linear contracts, then the purchaser can obtain the same welfare under a name and shame policy or a more sophisticated pay-for-performance scheme. More precisely, if the purchaser can implement a menu of contracts offering a different transfer in combination with a different price level: $\{T(\theta), p(\theta)\}$ with $\theta = \underline{\theta}, \overline{\theta}$, then allocative efficiency can be obtained for both doctors also under a separating equilibrium (proof available from the authors, and in line with Chone and Ma, 2011). Although these contracts are more flexible, we are not aware of policy examples which take such a form.

6 Equilibrium selection and the intuitive criterion test

In this section, we discuss two issues regarding the equilibria investigated in this paper. First, we discuss whether the beliefs that sustain the pooling equilibrium pass the Intuitive Criterion test (Cho and Kreps, 1987; Fudenberg and Tirole, 1991). Second, we address whether other Perfect Bayesian Equilibria coexist with the pooling equilibrium.

Pooling and the intuitive criterion. The pooling equilibrium ($q^E(\theta) = q^*(\bar{\theta})$ for all θ) is sustained by assuming that observing any quality *above* the one chosen by the good doctor $q^*(\bar{\theta})$ leads to the same beliefs (namely, $\theta^s = E(\theta)$) as observing the quality of the good doctor. Such (out-of-equilibrium) beliefs fail the Intuitive Criterion Test (ICT) if there exists a quality $q^h > q^*\left(\bar{\theta}\right)$ such that (i) the good doctor prefers quality q^h to $q^*\left(\bar{\theta}\right)$ if the former brings the best possible beliefs about altruism; and (ii) the bad doctor prefers his non-reputational optimum $q^*\left(\underline{\theta}\right)$ to q^h even if the latter leads to the best possible beliefs. If (i) and (ii) are satisfied at quality q^h then the Intuitive Criterion demands that society put zero probability on the doctor being bad upon observing such quality. This would contradict our assumption on the beliefs assigned to quality levels above the good doctor's quality. It turns out that a quality q^h satisfying conditions (i) and (ii) does exist for all reputation concerns $\alpha > \tau/\lambda$ if the patient's benefit monotonically increases in quality so that the single-crossing condition $V_{\theta q} \equiv W_q\left(q\right) > 0$ is satisfied for all q. In contrast, if the marginal benefit from quality is negative for high intensity of care this is not necessarily so. We summarize our results in the next proposition (proofs in Appendix 5).

Proposition 5 *Suppose that the benefit function is quadratic, $W(q) = v_1 q - v_2 q^2$, so that the marginal benefit from quality (intensity of care) can be negative, and cost is linear, $C(q) = cq$, where v_1, v_2, c are positive parameters, and that the purchaser sets price to achieve allocative efficiency, $p^P = c\left(1 - \bar{\theta}\right)$. Then there exists a threshold for reputational concerns $\alpha > \alpha^* \overset{def}{=} \frac{\bar{\theta}}{\bar{\theta}-\underline{\theta}} \frac{c^2}{v_2} \frac{1}{1-\lambda}$, such that the pooling equilibrium passes the Intuitive Criterion Test.*

The proposition is illustrated in Figure 3 (drawn, without loss of generality, for $\underline{\theta} = \frac{1}{4}$, $\bar{\theta} = \frac{3}{4}$, $c = 1$, $v_1 = 5$ and $v_2 = \frac{1}{8}$). The horizontal line represents the threshold of reputational concerns α below which the separating equilibrium arises, τ. The hyperbola represents the threshold of reputational concerns above which the pooling equilibrium arises, τ/λ. For reputational concerns between the hyperbola and the lower horizontal line the semiseparating equilibrium arises. The increasing curve depicts the threshold for reputational concerns α^* such that the pooling equilibrium passes the intuitive criterion test. The threshold α^* decreases in the good doctor's altruism and increases in the bad doctor's altruism. Therefore, the pooling equilibrium is intuitive if the good doctor has high altruism and if the bad doctor has low altruism. If the good doctor has high altruism, he is very sensitive to reductions in patient's benefits, which will occur if his care is excessive. If society observes such an intense treatment then it becomes implausible that it comes from a good doctor, and there does not exist any intensity of treatment such that only good doctors would choose. Similarly, if the bad doctor has low altruism, he tends to ignore the fact that patients' marginal benefits are negative when excessive care is provided. This implies that if society observes such excessive care then it becomes more plausible that it comes from a bad doctor.

[Figure 3 here]

Other equilibria and the Riley Outcome. As it is usually the case for signalling games, many PBEs coexist in large regions of parameter values. We have restricted attention to equilibria where (i) at least the good doctor sets his non-reputational optima and (ii) monotonicity and pessimism are imposed on out-of-equilibrium beliefs. We derive another PBE *candidate* that the literature has often focused on: the so-called Riley Outcome (Riley, 1979). In this (separating) outcome, the good doctor sets some (large) quality $\tilde{q}$ and the bad doctor sets his non-reputational optimum quality $q^*(\underline{\theta})$. The aforementioned quality $\tilde{q}$ is such that the bad doctor is indifferent between $\tilde{q}$ and his non-reputational optimum $q^*(\underline{\theta})$. It is also well known that if the single-crossing condition holds ($W(q)$ is monotone), then the Riley Outcome is a PBE, and moreover is the only such equilibrium that passes the ICT. In contrast, we show that under the assumptions listed in the previous proposition, the Riley Outcome ceases to be a PBE. We summarize the results in the next proposition.

Proposition 6 *If the marginal benefit from quality (intensity of care) can be negative, the Riley Outcome is not a PBE under the same assumptions as in Proposition 5 when reputational concerns are high,* $\alpha > \alpha^{RO} \overset{def}{=} \frac{\bar{\theta}}{\bar{\theta}-\underline{\theta}} \frac{c^2}{v_2}\left(\frac{1}{2}\sqrt{\frac{\bar{\theta}}{\underline{\theta}}} + \frac{\bar{\theta}+\underline{\theta}}{4\underline{\theta}}\right)$.

When reputational concerns are high, the bad doctor has an incentive to mimic the good doctor even when this requires high intensity of treatment. This is reinforced by the fact that the bad doctor is relatively insensitive to a decrease in patient's benefit due to excessive treatment. Therefore, the minimal intensity of treatment ($\tilde{q}$) that avoids imitation is far from (and above of) the good doctor's non-reputational optimum, where marginal benefits are already negative. For high reputational concerns (ie $\alpha > \alpha^{RO}$), intensity of treatment $\tilde{q}$ entails such a low benefit for the patient that the good doctor prefers his non-reputational optimum, where patients' benefit is higher, even if this brings a reputational loss. Hence the good doctor deviates from his Riley-Outcome strategy. Figure 4 illustrates (drawn for the same parameter values as Figure 3). This figure includes the (upper) horizontal line representing the threshold such that, for reputational concerns $\alpha > \alpha^{RO}$, the Riley Outcome is not a PBE. The most favorable region for supporting the pooling equilibrium is the one above line α^{RO} and enclosed between curves τ/λ and α^*. Indeed, in this region only the pooling equilibrium is a PBE and moreover it passes the ICT. This region requires sufficiently high reputational concerns and intermediate values on the proportion of good doctors.

[Figure 4 here]

7 Extensions

7.1 Risk averse doctors

In this section, we expand the analysis to allow doctors to be risk averse. A doctor's non-reputational payoff is in this case given by $\widetilde{V}(q) = u(T + pq) + \theta W(q) - C(q)$, with $u' > 0$ and $u'' \leq 0$, which is still concave in quality and increasing in altruism. The first order condition for quality $q^*(\theta)$ becomes $pu'(T + pq) + \theta W_q(q) = C_q(q)$. The thresholds for the three equilibrium regimes, τ and τ/λ, are given by the expressions in Section 2. The only difference is that $V(q)$ is replaced with $\widetilde{V}(q)$ so that $\tau = \widetilde{V}(q^*(\underline{\theta}) \mid \underline{\theta}) - \widetilde{V}(q^*(\overline{\theta}) \mid \underline{\theta}))/(\overline{\theta} - \underline{\theta})$. Our key results are therefore qualitatively similar under risk aversion. In Section 3, we have shown that the effect of an increase in price p per unit of quality is generally ambiguous under risk neutrality. This is still the case under risk aversion, but we can show that if the fixed payment T is relatively small, then crowding in is reinforced if the index of relative risk aversion is less than one and the index of absolute risk aversion is increasing ($u''' \leq 0$). There is neither crowding in nor crowding out under constant relative risk aversion (eg if the utility function is logarithmic in consumption). (Proofs available from the authors and omitted for brevity.)

7.2 Overprovision of quality

The main model assumes that quality is under-provided and that a pooling equilibrium is always an improvement since it induces the bad doctor to provide higher quality. In this section, we discuss how the model can be adapted to the case of over-provision of quality. Suppose that doctors are reimbursed a proportion of the cost γ,[8] and their profit function is $\widetilde{\pi}(q) = T + pq - (1-\gamma)C(q)$. The optimal quality $q^*(\theta)$, which maximises the non-reputational payoff, satisfies $p + \theta W_q = (1-\gamma)C_q$. In line with Section 5, suppose that the first-best quality for the purchaser q^o is such that it induces allocative efficiency, $W_q(q^o) = C_q(q^o)$. Then, in the absence of reputational concerns, it is straightforward to show that if altruism is less than one, over-provision of quality of the doctor with high altruism arises for sufficiently high degree of cost reimbursement, $q^*(\overline{\theta}) > q^o$ (this could also hold for the doctor with low altruism, for very high levels of cost reimbursement, which we rule out). The key insights of the main model still hold: the thresholds for the three equilibrium regimes, τ and τ/λ, are given by the same expressions as in Section 2. The only difference is that $\pi(q)$ is replaced with $\widetilde{\pi}(q)$. When reputational concerns are high, the doctor with low altruism mimics the doctor with high altruism and provides a higher quality, which is above the first-best level. In other words, name and shame policies now

[8] Such a reimbursement system falls in the broad category of supply-side cost sharing. For instance, most Health Maintenance Organizations in the US compensate their doctors through a flat salary plus a variable payment based on the volume of services they provide (Ellis and McGuire, 1993).

induce over-provision of quality for both doctors. Instead, when reputational concerns are low, eg in the absence of name and shame policies, the doctor with low altruism under-provides quality and the doctor with high altruism over-provides quality. The welfare implications are ambiguous. Name and shame policies can reduce welfare if the welfare loss from over-provision of quality of the doctor with low altruism under pooling is higher than the welfare loss of under-provision of quality from the doctor with low altruism in the separating equilibrium. The same would hold under a multitasking set-up as in Section 4 so that over-provision of observed quality can co-exist with under-provision of unobserved quality.

7.3 Reputational concerns and doctors' competition

The main model assumes that reputational concerns for doctors are non-monetary and are determined by what other people (their family, the community, their peers, other patients, and more broadly society) think about them. An alternative interpretation is that reputational concerns arise from future monetary payoffs if, for example, a better reputation brings future patient demand. This would require a dynamic set-up, in which doctors care about their reputation in Period 1 due to non-monetary motives but also because they compete for demand in Period 2. Within a model *a la Hotelling*,[9] we can show that the pooling, separating and semi-separating equilibrium obtained in Section 2 also arise as a function of the *sum* of the non-monetary (as in Section 2) and the monetary reputational concerns from future demand (proof omitted but available from the authors).

8 Conclusions

The health sector has witnessed a proliferation of performance indicators in the public domain. Can the mere publishing of information on the quality of doctors induce them to change behaviour and work harder? The analysis of this study suggests *Yes, it can.* Policies colloquially known as *name and shame*, in which poorly performing doctors are subjected to shame in front of the community, can induce the poor performing doctors to provide more effort to avoid being tagged as bad doctors, a form of virtuous imitation.

[9] Suppose that in period 1 there is a continuum of patients uniformly distributed on a unit line, with two doctors, A and B, located in the extremes of the unit line. Patients are uninformed about whether doctors are good or bad and therefore visit the nearest doctor. Hence each doctor treats half of the market in period 1, and the share of the market in period 1 does not convey any information about doctor's altruism. Patients are homogeneous in benefits and treatment costs, and these are given by $W(q)$ and $C(q)$ as in Section 2. Each doctor $i = A, B$ chooses quality q_i to treat her patients in period 1. Once these patients are treated, they are cured and disappear from the market. In period 2, a *new* set of patients appears. These patients are again uniformly distributed on the unit line. The net revenue for each patient treated is $m \geq 0$ which is borne by a third party payer. Patients in period 2 choose doctors based on the posterior beliefs about doctors' altruism. These patients care about these beliefs because their benefit from visiting a doctor increases with this doctor's altruism.

Moreover, we have shown that pay-for-performance schemes are not a perfect substitute for policies which disseminate information. Publishing indicators can raise quality even if incentive schemes are optimally set by purchasers, as long as the purchaser is constrained to adopt relatively simple contracts (ie in the absence of menus of contracts, which are rarely observed). Our results are good news also in terms of equity. The presence of sufficiently strong reputational concerns always reduces the gap between the quality of the good and the bad doctor.

Our model assumes that doctors have perfect information on the benefits and costs of quality. It may be argued that for some treatments they may have imperfect information about the appropriate quality. Under such scenarios doctors may respond to quality reports if they convey information about the appropriate care. We leave this issue for future research.

9 References

Andreoni, J., 1989. Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence. Journal of Political Economy 97 (6), 1447-1458.

Ariely, D., Bracha, A., Meier, M., 2009. Doing Good or Doing Well? Image Motivation and Monetary Incentives in Behaving Prosocially. American Economic Review, 99 (1), 544-55.

Bénabou, R., Tirole, J., 2006. Incentives and prosocial behavior. American Economic Review 96 (5), 1652–1678.

Bénabou, R., Tirole, J., 2011. Laws and Norms. NBER Working paper N. 17579, 1-44.

Besley, T., Ghatak, M., 2005. Competition and incentives with motivated agents. American Economic Review 95 (3), 616–636.

Brekke, K., Siciliani, L., Straume, O.R., 2011. Hospital Competition and Quality with Regulated Prices. Scandinavian Journal of Economics 113, 444-469.

Brekke, K., Siciliani, L., Straume, O.R., 2012. Quality competition with profit constraints. Journal of Economic Behaviour & Organization 84, 642-659.

Cartwright, E., Patel, A., 2013. How category reporting can improve fundraising. Journal of Economic Behavior & Organization 87, 73–90.

Cashin, C., Chi, Y.-L., Smith, P., Borowitz, M., Thomson, S., 2014. Paying for Performance in Healthcare: Implications for Health System Performance and Accountability. Open University Press, Buckingham, UK.

Chalkley, M., Malcomson, J.M., 1998. Contracting for health services when patient demand does not reflect quality. Journal of Health Economics 17, 1–19.

Chen, Y., 2011. Why are health care report cards so bad (good)? Journal of Health Economics, 30, 575-590.

Cho, I., Kreps, D., 1987. Signalling games and stable equilibria. Quarterly Journal of Economics 102 (2), 179–221.

Choné, P., Ma, C.A., 2011. Optimal health care contracts under physician agency. Annales d'Economie et de Statistique 101/102, 229-256.

Daugherty, A.F., Reinganum, J.F., 2010. Public Goods, Social Pressure, and the Choice between Privacy and Publicity. American Economic Journal: Microeconomics, American Economic Association 2 (2), 191-221.

Delfgaauw, J., Dur, R., 2008. Incentives and Workers. Motivation in the Public Sector. Economic Journal 118, 171-191.

Dixit, A., 2005. Incentive contracts for faith-based organisations to deliver social services. In: Lahiri, Sajal, Maiti, Pradip (Eds.), Economic Theory in a Changing World: Policy Modelling for Growth. Oxford University Press, New Delhi.

Doran, T. et al., 2006. Pay-for-Performance Programs in Family Practices in the United Kingdom. New England Journal of Medicine 355 (4), 375-384.

Dranove, D., Kessler, D., McClellan, M. Satterthwaite, M., 2003. Is more information better? The effects of 'report cards' on health care providers, Journal of Political Economy 111, 555-558.

Eggleston K. 2005. Multitasking and mixed systems for provider payment. Journal of Health Economics 24 (1), 211-223.

Ellis, R.P., McGuire, T.G., 1986. Provider Behavior under Prospective Reimbursement. Journal of Health Economics 5, 129-51.

Ellis, R.P., McGuire, T.G., 1993. Supply-Side and Demand-Side Cost Sharing in Health Care. The Journal of Economic Perspectives 7, 135-151.

Falk, A., Kosfeld, M., 2006. The hidden costs of control. American Economic Review 96 (5), 1611-1630.

Fehr, E., List, J.A., 2004. The Hidden Costs and Returns of Incentives-Trust and Trustworthiness Among CEOs. Journal of the European Economic Association, 2 (5), 743-771.

Fichera, E., Nikolova, S., and Sutton, M., 2014. Comparative Performance Evaluation: Quality, Elsevier Encyclopaedia of Health Economics, ed. T. Culyer, 111-116.

Francois, P., 2000. 'Public service motivation' as an argument for government provision. Journal of Public Economics 78, 275-299.

Fudenberg, D., Tirole, J., 1991. Game theory, MIT Press.

Funk, P. 2007. Is There An Expressive Function of Law? An Empirical Analysis of Voting Laws with Symbolic Fines, American Law and Economics Review 9, 135-159.

Glazer, A., 2004. Motivating devoted workers. International Journal of Industrial Organization 22, 427– 440.

Godager, G., Wiesen, D., 2013. Profit or patients' health benefit? Exploring the heterogeneity in physician altruism, Journal of Health Economics 32, 1105-1116.

Gravelle, H., Sutton, M., Ma, A., 2010. Doctor Behaviour under a Pay for Performance Contract: Treating, Cheating and Case Finding?, Economic Journal 120, F129-F56.

Hibbard, J.H., Stockard, J., and Tusler, M., 2005. Hospital Performance Reports: Impact On Quality, Market Share, And Reputation, Health Affairs 24, 1150-1160.

Holmstrom, B., Milgrom, P., 1991. Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design, Journal of Law, Economics and Organization 7, 24-52.

Jack, W., 2005. Purchasing health care services from providers with unknown altruism. Journal of Health Economics 24, 73-93.

Jeitschko, T.D. and Normann, H.-T., 2012. Signaling in deterministic and stochastic settings. Journal of Economic Behavior & Organization 82, 39– 55.

Kaarboe, O., Siciliani, L., 2011. Quality, Multitasking and Pay for Performance. Health Economics 2, 225-238.

Kolstad, J.T., 2013. Information and Quality When Motivation is Intrinsic: Evidence from Surgeon Report Cards. American Economic Review 103 (7), 2875–2910.

Lakdawalla, D., Philipson, T., 2006. The nonprofit sector and industry performance. Journal of Public Economics 90, 1681–1698.

Lindenauer, P.K., Remus, D. Roman, S., Rothberg, M.B., Benjamin, E.M., Ma, A. Bratzler, D.W., 2007. Public Reporting and Pay for Performance in Hospital Quality Improvement, The New England Journal of Medicine, 356, 486-496.

Lu, S. F., 2012. Multitasking, information disclosure, and product quality: Evidence from nursing homes. Journal of Economics & Management Strategy 21; 673–705.

Makris, M., 2009. Incentives for Motivated Agents under an Administrative Constraint. Journal of Economic Bahavior and Organization 71, 428-440.

Makris, M., Siciliani, L., 2013. Optimal Incentive Schemes for Altruistic Providers. Journal of Public Economic Theory 15 (5), 675–699.

Prendergast, C., 2007. The motivation and bias of bureaucrats. American Economic Review 97, 180–96.

Riley, J.G. 1979. Informational Equilibrium. Econometrica 47 (2), 331-359.

Rodriguez-Barraquer, T., Xu, T. 2015. A Model of Competitive Signaling, mimeo, Universitat Autònoma de Barcelona.

Roland, M., Dudley, R.A., 2015. How financial and reputational incentives can be used to improve medical care. Health Service Research 50, 2090-2115.

Rosenthal, M.B., France, R.G., Li Z., Epstein, A.M., 2005. Early Experience with Pay-for-Performance: From Concept to Practice, JAMA, The Journal of the American Medical Association 294 (14), 1788-1793.

Rosenthal, M.B., Fernandopulle, R., Song, H.R., Landon, B., 2004. Paying for quality: providers' incentives for quality improvement. Health Affairs 23 (2), 127-41.

Ryan, R. M., Deci, E.L., 2000. Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being, American Psychologist 55, 68-78.

Siciliani, L., 2009. Pay for performance and motivation crowding out. Economics Letters 103, 68–71.

Siciliani, L., Straume, O.R., Cellini, R., 2013. Quality competition with motivated providers and sluggish demand, Journal of Economic Dynamics and Control 37, 2041-2061.

Wang, J., Hockenberry, J., Chou, S., Yang, M., 2011. Do bad report cards have consequences? Impacts of publicly reported provider quality information on the CABG market in Pennsylvania, Journal of Health Economics 30 (2), 392-407.

Werner, R. M., Konetzka, R. T., Kruse, G. B., 2009. Impact of public reporting on unreported quality of care, Health Services Research 44, 379–398.

Appendix

Appendix 1. The model. Our equilibrium concept to solve the model is the Perfect Bayesian Equilibrium: An equilibrium is a pair of functions of qualities $q^E(\theta) : \{\underline{\theta}; \overline{\theta}\} \rightarrow R_+$ and beliefs $\lambda^S(q) : R_+ \rightarrow [0,1]$ such that (i) for every θ in $\{\underline{\theta}; \overline{\theta}\}$, $q^E(\theta)$ maximizes $\pi(q)+\theta W(q)+\alpha(\theta^s(q)-E(\theta))$ with $\theta^S(q) = \lambda^S(q)\bar{\theta}+\left(1-\lambda^S(q)\right)\underline{\theta}$, (ii) $\lambda^S(q)$ is computed using Bayes' rule whenever possible, and (iii) $\lambda^S(q)$ is any number between 0 and 1 when Bayes' rule cannot be applied. Bayes' rule cannot be applied when the observed quality q is neither types' posited equilibrium choice (ie the denominator of Bayes' formula is zero). As an example of an out of equilibrium action, suppose that in (a pooling) equilibrium all types set $q = q_0$, then any $q \neq q_0$ is a non-equilibrium action.

Separating equilibrium. Out of equilibrium beliefs should satisfy monotonicity and pessimism, so $\lambda^S(q) = 0$ (so $\theta^s(q) = \underline{\theta}$) for any $q < q^*(\underline{\theta})$ and $\lambda^S(q) = 1$ (so $\theta^s(q) = \bar{\theta}$) for any $q \geq q^*(\underline{\theta})$. For these strategies and beliefs to constitute a separating PBE, we need the incentive-compatibility (IC) constraints for both types to be satisfied. The good doctor has no incentive to deviate to any other quality if

$$V(q^*(\bar{\theta}) \mid \bar{\theta}) + \alpha(\bar{\theta} - E(\theta)) \geq \left\{ \begin{array}{c} V(q \mid \bar{\theta}) - \alpha(E(\theta) - \underline{\theta}) \text{ for all } q < q^*\left(\bar{\theta}\right), \\ V(q \mid \bar{\theta}) + \alpha(\bar{\theta} - E(\theta)) \text{ for all } q \geq q^*\left(\bar{\theta}\right) \end{array} \right\},$$

and the bad doctor has no incentive to deviate to any other quality if

$$V(q^*(\underline{\theta}) \mid \underline{\theta}) - \alpha(E(\theta) - \underline{\theta}) \geq \left\{ \begin{array}{c} V(q \mid \underline{\theta}) - \alpha(E(\theta) - \underline{\theta}) \text{ if } q < q^*\left(\bar{\theta}\right), \\ V(q \mid \underline{\theta}) + \alpha(\bar{\theta} - E(\theta)) \text{ if } q \geq q^*\left(\bar{\theta}\right) \end{array} \right\}.$$

The IC constraint for the good doctor is always satisfied since the high quality $q^*\left(\bar{\theta}\right)$ maximizes the non-reputational payoff $V(q \mid \overline{\theta})$, and choosing this quality instead of any other quality also maximises the reputational payoff. As for the bad doctor, since the low quality $q^*(\underline{\theta})$ maximizes the non-reputational payoff $V(q \mid \underline{\theta})$, choosing such quality generates the same reputation loss as any other quality which is below the high quality $q^*\left(\bar{\theta}\right)$. The bad doctor must also be better-off by providing the low quality $q^*(\underline{\theta})$ rather than by disguising himself by providing the higher quality $q^*\left(\bar{\theta}\right)$ in the attempt of gaining the reputational payoff. Note that the bad doctor has no incentive to choose a quality which is strictly above the high quality $q^*\left(\bar{\theta}\right)$ since it would increase costs with no additional gains (ie $V(q^*\left(\bar{\theta}\right) \mid \underline{\theta}) + \alpha\left(\bar{\theta} - \underline{\theta}\right)$ is maximized at $q = q^*\left(\bar{\theta}\right)$ conditional on $q \geq q^*\left(\bar{\theta}\right)$). Hence the lower expression in the RHS of the bad doctor's IC implies $V(q^*(\underline{\theta}) \mid \underline{\theta}) - \alpha(E(\theta) - \underline{\theta}) \geq V(q^*\left(\bar{\theta}\right) \mid \underline{\theta}) + \alpha(\bar{\theta} - E(\theta))$.

Pooling equilibrium. According to our beliefs restrictions, we have that any $q < q^*(\bar{\theta})$ indicates that the doctor is bad, ie $\lambda^S(q) = 0$ (so $\theta^s(q) = \underline{\theta}$) for any $q < q^*\left(\bar{\theta}\right)$; and that any higher quality than $q^*\left(\bar{\theta}\right)$ does not provide any further information, ie

$\lambda^S(q) = \lambda$ (so $\theta^s(q) = E(\theta)$) for $q \geq q^*\left(\bar{\theta}\right)$. For these strategies and beliefs to constitute an equilibrium, we need again the incentive-compatibility constraints to be satisfied, that is, both types of doctor must have an incentive to provide the high quality, namely, $V(q^*(\bar{\theta}) \mid \bar{\theta}) \geq \left\{ \begin{array}{c} V(q \mid \bar{\theta}) \text{ for all } q \geq q^*\left(\bar{\theta}\right), \\ V(q \mid \bar{\theta}) - \alpha(E(\theta) - \underline{\theta}) \text{ for all } q < q^*\left(\bar{\theta}\right) \end{array} \right\}$ and $V(q^*(\bar{\theta}) \mid \underline{\theta}) \geq \left\{ \begin{array}{c} V(q \mid \underline{\theta}) \text{ for all } q \geq q^*\left(\bar{\theta}\right), \\ V(q \mid \underline{\theta}) - \alpha(E(\theta) - \underline{\theta}) \text{ for all } q < q^*\left(\bar{\theta}\right) \end{array} \right\}$. The good doctor's IC is always satisfied since (i) the non-reputational payoff $V(q \mid \bar{\theta})$ is maximized at the high quality $q^*(\bar{\theta})$ and, (ii) any other quality below $q^*(\bar{\theta})$ brings a reputational loss (equal to $\alpha\left(E(\theta) - \underline{\theta}\right)$). The bad doctor's IC is satisfied only if reputational concerns are sufficiently high. To see this, consider first the upper condition in his IC. Since the non-reputational payoff $V(q \mid \underline{\theta})$ is maximized at the low quality $q^*\left(\underline{\theta}\right) < q^*\left(\bar{\theta}\right)$, it decreases for any q in excess of $q^*\left(\bar{\theta}\right)$. Therefore, the condition reduces to $V(q^*(\bar{\theta}) \mid \underline{\theta}) \geq V(q^*\left(\underline{\theta}\right) \mid \underline{\theta}) - \alpha(E(\theta) - \underline{\theta})$, which can be re-written as $\alpha \geq \tau/\lambda$.

The semi-separating equilibrium. Assume $\alpha \in \left(\tau, \frac{\tau}{\lambda}\right)$. Suppose that the good doctor chooses the high quality $q^E\left(\bar{\theta}\right) = q^*\left(\bar{\theta}\right)$ with certainty, and the bad doctor chooses the low quality with probability r and the high quality with probability $(1-r)$. Then the equilibrium is characterised by $r = r^E \overset{\text{def}}{=} 1 - \frac{\lambda}{(1-\lambda)}\left(\frac{\alpha-\tau}{\tau}\right) < 1$; $\lambda^S\left(q^*\left(\underline{\theta}\right)\right) = 0$, $\lambda^S\left(q^*\left(\bar{\theta}\right)\right) = \frac{\tau}{\alpha} > \lambda$, $\theta^s\left(q^*\left(\underline{\theta}\right)\right) = \underline{\theta}$, and $\theta^S\left(q^*\left(\underline{\theta}\right)\right) = \underline{\theta} + \frac{\tau}{\alpha}\left(\bar{\theta} - \underline{\theta}\right) > E(\theta)$. Given the posited strategies, Bayes' Rule can always be applied to $q \in \left\{q^*\left(\bar{\theta}\right), q^*\left(\underline{\theta}\right)\right\}$. Posterior beliefs when either of these two qualities is observed are $\lambda^S\left(q^*\left(\underline{\theta}\right)\right) = 0$ and $\lambda^S\left(q^*\left(\bar{\theta}\right)\right) = \frac{\lambda}{1-r(1-\lambda)} > \lambda$. These beliefs yield the following expected types: $\theta^S\left(q^*\left(\underline{\theta}\right)\right) = \underline{\theta}$ and $\theta^S\left(q^*\left(\bar{\theta}\right)\right) = \frac{\lambda\bar{\theta}+(1-r)(1-\lambda)\underline{\theta}}{1-r(1-\lambda)}$. Any quality $q \notin \left\{q^*\left(\bar{\theta}\right), q^*\left(\underline{\theta}\right)\right\}$ is out of equilibrium. According to our restriction on out-of-equilibrium beliefs, we have that $\lambda^S(q) = \left\{ \begin{array}{c} 0 \text{ for all } q < q^*\left(\bar{\theta}\right), \\ \frac{\lambda}{1-r(1-\lambda)} \text{ for all } q \geq q^*\left(\bar{\theta}\right) \end{array} \right\}$. Hence the expected type upon observation of such q is $\theta^S(q) = \left\{ \begin{array}{c} \underline{\theta} \text{ for all } q < q^*\left(\bar{\theta}\right), \\ \frac{\lambda}{1-r(1-\lambda)}\bar{\theta} + \frac{(1-\lambda)(1-r)}{1-r(1-\lambda)}\underline{\theta} \text{ for all } q \geq q^*\left(\bar{\theta}\right) \end{array} \right\}$. We can now determine the reputational payoff: $G(q) = \left\{ \begin{array}{c} -\alpha\left(E(\theta) - \underline{\theta}\right) \text{ if } q < q^*\left(\bar{\theta}\right), \\ \alpha\left(\frac{\lambda\bar{\theta}+(1-r)(1-\lambda)\underline{\theta}}{1-r(1-\lambda)} - E(\theta)\right) \text{ if } q \geq q^*\left(\bar{\theta}\right) \end{array} \right\}$. For these strategies and beliefs to constitute a PBE we need three conditions. First, the bad doctor has to be indifferent between $q = q^*\left(\bar{\theta}\right)$ and $q = q^*\left(\underline{\theta}\right)$; second, the bad doctor has to (weakly) prefer any of the latter to setting $q \notin \left\{q^*\left(\bar{\theta}\right), q^*\left(\underline{\theta}\right)\right\}$; and third, the good doctor has to weakly prefer $q = q^*\left(\bar{\theta}\right)$ to $q \neq q^*\left(\bar{\theta}\right)$ despite the fact that a high quality does not fully reveal his type. Using the fact that $q^*(\cdot)$ maximizes $V(q|\cdot)$, and that $\frac{\lambda\bar{\theta}+(1-r)(1-\lambda)\underline{\theta}}{1-r(1-\lambda)} - \underline{\theta} = \lambda\frac{\bar{\theta}-\underline{\theta}}{1-r(1-\lambda)}$, these three conditions can be written as

$$V(q^*(\bar{\theta}) \mid \underline{\theta}) + \alpha\lambda \frac{\bar{\theta} - \underline{\theta}}{1 - r(1 - \lambda)} = V(q^*(\underline{\theta}) \mid \underline{\theta}), \tag{10}$$

$$V(q^*(\underline{\theta}) \mid \underline{\theta}) \geq \begin{cases} V(q \mid \underline{\theta}) \text{ for all } q < q^*(\bar{\theta}) \\ V(q \mid \underline{\theta}) + \alpha\left(\lambda \frac{\bar{\theta} - \underline{\theta}}{1 - r(1 - \lambda)}\right) \text{ for } q > q^*(\bar{\theta}), \end{cases} \tag{11}$$

$$V(q^*(\bar{\theta}) \mid \bar{\theta}) \geq \begin{cases} V(q \mid \bar{\theta}) - \alpha\lambda \frac{\bar{\theta} - \underline{\theta}}{1 - r(1 - \lambda)} \text{ for all } q < q^*(\bar{\theta}) \\ V(q \mid \bar{\theta}) \text{ for all } q \geq q^*(\bar{\theta}). \end{cases} \tag{12}$$

Notice that the upper expression in (11) is always satisfied because $q^*(\underline{\theta})$ maximises $V(q \mid \underline{\theta})$. The lower expression in (11) can be rewritten, using (10), as $V(q^*(\bar{\theta}) \mid \underline{\theta}) \geq V(q \mid \underline{\theta})$ for $q > q^*(\bar{\theta})$. This condition is again satisfied since $q^*(\underline{\theta})$ maximises $V(q \mid \underline{\theta})$. Similarly, (12) is also always satisfied since $V(q \mid \bar{\theta})$ is largest at $q^*(\bar{\theta})$ and since $\alpha\lambda \frac{\bar{\theta} - \underline{\theta}}{1 - r(1 - \lambda)} > 0$ because $\lambda, r \in (0, 1)$. Hence only (10) is restrictive and equivalent to $r = 1 - \frac{\lambda}{(1 - \lambda)} \left(\frac{\alpha - \tau}{\tau}\right) = r^E$. Substituting this expression into the expression for $\lambda^S\left(q^*(\bar{\theta})\right)$, we obtain (after some algebra) $\lambda^S\left(q^*(\bar{\theta})\right) = \frac{\tau}{\alpha}$. Then $\theta^S\left(q^*(\underline{\theta})\right) = \underline{\theta} + \frac{\tau}{\alpha}\left(\bar{\theta} - \underline{\theta}\right)$.

Equilibrium payoffs. Let us use $\theta^S\left(q^*(\underline{\theta})\right) = \underline{\theta} + \frac{\tau}{\alpha}\left(\bar{\theta} - \underline{\theta}\right)$. To sustain a mixed strategy with support $\{q^*(\bar{\theta}), q^*(\underline{\theta})\}$, the bad doctor must be indifferent between these two qualities. Hence his payoff is given by $V(q^*(\underline{\theta}) \mid \underline{\theta}) - \alpha(E(\theta) - \underline{\theta}) = V(q^*(\bar{\theta}) \mid \underline{\theta}) + (\tau - \alpha\lambda)\left(\bar{\theta} - \underline{\theta}\right)$. The good doctor obtains $V(q^*(\bar{\theta}) \mid \bar{\theta}) + \alpha\left[\theta^S\left(q^*(\bar{\theta})\right) - E(\theta)\right] = V(q^*(\bar{\theta}) \mid \bar{\theta}) + (\tau - \alpha\lambda)\left(\bar{\theta} - \underline{\theta}\right)$. Notice that the good doctor's payoff tends to $V(q^*(\bar{\theta}) \mid \bar{\theta}) + \tau(1 - \lambda)\left(\bar{\theta} - \underline{\theta}\right)$ when α tends to τ, which is the same as the separating payoff at $\alpha = \tau$, and that it tends to $V(q^*(\bar{\theta}) \mid \bar{\theta})$ when α tends to τ/α, which is the same as the pooling payoff. This is an interesting feature of the semi-separating equilibrium: it connects the separating and the pooling equilibrium. This can be easily checked by inspection.

Appendix 2. Crowding in and crowding out. Using the envelop theorem, we have that: $\frac{\partial V(q^*(\underline{\theta})|\underline{\theta})}{\partial p} = q^*(\underline{\theta}) > 0$: a higher price increases revenues and therefore the utility of the bad doctor when the optimal quality $q^*(\underline{\theta})$ is chosen. In contrast

$$\frac{\partial V(q^*(\bar{\theta}), \underline{\theta})}{\partial p} = q^*(\bar{\theta}) + [p + \underline{\theta} W_q(q^*(\bar{\theta})) - C_q(q^*(\bar{\theta}))] \frac{\partial q^*(\bar{\theta})}{\partial p} \gtreqless 0. \tag{13}$$

Higher prices increase revenues but also increase the quality of the good doctor, which makes it more costly for the bad doctor to disguise as the good doctor. By substitution, we therefore obtain $\frac{\partial \tau}{\partial p} = \frac{q^*(\bar{\theta}) - q^*(\underline{\theta})}{\bar{\theta} - \underline{\theta}} - \frac{p - C_q(q^*(\bar{\theta})) + \underline{\theta} W_q(q^*(\bar{\theta}))}{\bar{\theta} - \underline{\theta}} \frac{\partial q^*(\bar{\theta})}{\partial p}$. From the FOC of quality of the good doctor we have $p - C_q(q^*(\bar{\theta})) = -\bar{\theta} W_q(q^*(\bar{\theta}))$ which we substitute in $\frac{\partial \tau}{\partial p}$. The result is obtained: $\frac{\partial \tau}{\partial p} = -\frac{q^*(\bar{\theta}) - q^*(\underline{\theta})}{\bar{\theta} - \underline{\theta}} + W_q(q^*(\bar{\theta})) \frac{\partial q^*(\bar{\theta})}{\partial p}$. Notice that $\frac{\partial q^*(\bar{\theta})}{\partial \bar{\theta}} = \frac{W_q(q^*(\bar{\theta}))}{-\bar{\theta} W_{qq}(q^*(\bar{\theta})) + C_{qq}(q^*(\bar{\theta}))} = W_q(q^*(\bar{\theta})) \frac{\partial q^*(\bar{\theta})}{\partial p}$. By substitution we obtain $\frac{\partial \tau}{\partial p} = \frac{\partial q^*(\bar{\theta})}{\partial \bar{\theta}} - \frac{q^*(\bar{\theta}) - q^*(\underline{\theta})}{\bar{\theta} - \underline{\theta}}$. The effect of prices on τ then depends on the concavity or convexity of quality as a function of altruism $q^*(\theta)$. If the function is concave (convex), ie $\frac{\partial^2 q^*(\theta)}{\partial^2\theta} < (>)0$, then $\frac{\partial q^*(\bar{\theta})}{\partial\bar{\theta}} < (>)\frac{q^*(\bar{\theta})-q^*(\underline{\theta})}{\bar{\theta}-\underline{\theta}}$. Therefore $\frac{\partial\tau}{\partial p}$ has the same sign as

$$\frac{\partial^2 q^*(\theta)}{\partial^2\theta} = W_q \frac{2W_{qq} + (\theta W_{qqq} - C_{qqq})\,(\partial q^*(\theta)/\partial\theta)}{(-\theta W_{qq} + C_{qq})^2}.$$

A sufficient condition for crowding in is $W_{qq} < 0$, $W_{qqq} \leq 0$ and $C_{qqq} \geq 0$ (the latter is always satisfied if the cost is quadratic or linear since $C_{qqq} = 0$). For crowding out to arise, the marginal benefit from quality has to be constant or mildly decreasing and the marginal cost has to be concave in quality (for instance if $C(q) = \beta q^n$ for $1 < n < 2$).

Appendix 3. Multitasking. The problem is well behaved, and the SOCs are: $V_{q_1q_1}(q_1^*, q_2^*) = \theta W_{q_1q_1}(q_1^*) - C_{q_1q_1}(q_1^*, q_2^*) < 0$, $V_{q_2q_2}(q_1^*, q_2^*) = \theta W_{q_2q_2}(q_2^*) - C_{q_2q_2}(q_1^*, q_2^*) < 0$ and $V_{q_1q_1}V_{q_2q_2} - C_{q_1q_2}^2 > 0$. The effect of price on qualities is: $\frac{\partial q_1^*}{\partial p} = -\frac{V_{q_2q_2}}{V_{q_1q_1}V_{q_2q_2}-C_{q_1q_2}^2} > 0$, $\frac{\partial q_2^*}{\partial p} = -\frac{C_{q_1q_2}}{V_{q_1q_1}V_{q_2q_2}-C_{q_1q_2}^2} < 0$. The effect of altruism on quality is: $\frac{\partial q_1^*}{\partial\theta} = -\frac{W_{q_1}V_{q_2q_2}+W_{q_2}C_{q_1q_2}}{V_{q_1q_1}V_{q_2q_2}-C_{q_1q_2}^2}$, $\frac{\partial q_2^*}{\partial\theta} = -\frac{W_{q_2}V_{q_1q_1}+W_{q_1}C_{q_1q_2}}{V_{q_1q_1}V_{q_2q_2}-C_{q_1q_2}^2}$. An equilibrium is a pair of functions of qualities $q_1^E(\theta), q_2^E(\theta) : \{\underline{\theta};\bar{\theta}\} \rightarrow R_+$ and beliefs $\lambda^S(q_1) : R_+ \rightarrow [0,1]$ such that: (i) for every θ in $\{\underline{\theta};\bar{\theta}\}$, $q_1^E(\theta), q_2^E(\theta)$ maximises $\pi(q_1, q_2)+\theta W(q_1, q_2)+G(q_1)$ once $\theta^S(q_1) = \lambda^S(q_1)\bar{\theta}+\left(1 - \lambda^S(q_1)\right)\underline{\theta}$ has been substituted into $G(q_1)$, (ii) $\lambda^S(q_1)$ is computed using Bayes' rule whenever possible, and (iii) $\lambda^S(q_1)$ is any number between 0 and 1 when Bayes' rule cannot be applied.

Pooling equilibrium. Since doctors provide the same quality $q_1^E(\theta)$, patients and society cannot distinguish between good and bad doctors. There is therefore no updating in beliefs after observing the high quality $q_1^*(\bar{\theta})$. Hence $\lambda^S\left(q_1^*(\bar{\theta})\right) = \lambda$ and the expected type conditional on patients observing the high quality is the average type, $\theta^s\left(q_1^*(\bar{\theta})\right) = E(\theta)$. Moreover, according to our beliefs restrictions (again monotonicity and pessimism) we have that any smaller quality than $q_1^*(\bar{\theta})$ implies that the doctor is bad, ie $\lambda^S(q_1) = 0$ (so $\theta^s(q_1) = \underline{\theta}$) for any $q_1 < q_1^*(\bar{\theta})$; and that any higher quality than $q_1^*(\bar{\theta})$ does not provide any further information, ie $\lambda^S(q_1) = \lambda$ (so $\theta^s(q_1) = E(\theta)$) for $q_1 \geq q_1^*(\bar{\theta})$. For the posited strategies and beliefs to constitute a pooling equilibrium, we need the incentive-compatibility (IC) constraints to be satisfied, so that both types of doctor have an incentive to provide the high quality. The good doctor has no incentive deviate if

$$V(q_1^*(\bar{\theta}), q_2^*(\bar{\theta}) \mid \bar{\theta}) \geq \begin{cases} V(q_1, q_2 \mid \bar{\theta}) \text{ for all } q_1 \geq q_1^*(\bar{\theta}) \\ V(q_1, q_2 \mid \bar{\theta}) - \alpha(E(\theta) - \underline{\theta}) \text{ for all } q_1 < q_1^*(\bar{\theta}). \end{cases} \quad (14)$$

The bad doctor has no incentive to deviate if

$$V(q_1^*(\bar{\theta}), q_2^M(\underline{\theta}) \mid \underline{\theta}) \geq \begin{cases} V(q_1, q_2 \mid \underline{\theta}) \text{ for all } q_1 \geq q_1^*(\bar{\theta}) \\ V(q_1, q_2 \mid \underline{\theta}) - \alpha(E(\theta) - \underline{\theta}) \text{ for all } q_1 < q_1^*(\bar{\theta}). \end{cases} \quad (15)$$

The IC constraint for the good doctor is always satisfied since (i) the non-reputational payoff $V(q_1, q_2 \mid \bar{\theta})$ is maximized at the high qualities $q_1^*(\bar{\theta})$ and $q_2^*(\bar{\theta})$ and, (ii) any other quality below $q_1^*(\bar{\theta})$ brings a reputational loss (equal to $-\alpha\left(E\left(\theta\right) - \underline{\theta}\right)$). The IC constraint for the bad doctor is satisfied only if reputational concerns are sufficiently high. Since the non-reputational payoff $V(q_1, q_2 \mid \underline{\theta})$ is maximized at the low qualities $q_1^*(\underline{\theta})$ and $q_2^*(\underline{\theta})$, it decreases for any q_1 in excess of $q_1^*\left(\bar{\theta}\right)$. Therefore, the condition reduces to

$$V(q_1^*(\bar{\theta}), q_2^M\left(\underline{\theta}\right) \mid \underline{\theta}) \geq V(q_1^*(\underline{\theta}), q_2^*(\underline{\theta}) \mid \underline{\theta}) - \alpha(E\left(\theta\right) - \underline{\theta}), \tag{16}$$

which can be re-written as $\alpha \geq \frac{\tau_1}{\lambda} \overset{def}{=} \frac{V(q_1^*(\underline{\theta}), q_2^*(\underline{\theta}) \mid \underline{\theta}) - V(q_1^*(\bar{\theta}), q_2^M(\underline{\theta}) \mid \underline{\theta})}{\lambda.\left(\bar{\theta} - \underline{\theta}\right)}$.

Separating equilibrium. Beliefs in the equilibrium path are such that observing a high (low) quality signals with certainty high (low) altruism: $\lambda^S\left(q_1^*\left(\bar{\theta}\right)\right) = 1$ and $\lambda^S\left(q_1^*\left(\underline{\theta}\right)\right) = 0$, where recall that only q_1 is observable. Out of equilibrium beliefs should satisfy monotonicity and pessimism, so $\lambda^S\left(q_1\right) = 0$ (so $\theta^s\left(q_1\right) = \underline{\theta}$) for any $q < q_1^*\left(\underline{\theta}\right)$ and $\lambda^S\left(q_1\right) = 1$ (so $\theta^s\left(q_1\right) = \bar{\theta}$) for any $q_1 \geq q_1^*\left(\underline{\theta}\right)$. For the posited strategies and beliefs to constitute a separating Bayesian Equilibrium, we need again the IC constraints for both types to be satisfied. The good doctor has no incentive to mimic the bad doctor:

$$V(q_1^*(\bar{\theta}), q_2^*(\bar{\theta}) \mid \bar{\theta}) + \alpha(\bar{\theta} - E\left(\theta\right)) \geq \begin{cases} V(q_1, q_2 \mid \bar{\theta}) - \alpha(E\left(\theta\right) - \underline{\theta}) \text{ for all } q_1 < q_1^*\left(\bar{\theta}\right) \\ V(q_1, q_2 \mid \bar{\theta}) + \alpha(\bar{\theta} - E\left(\theta\right)) \text{ for all } q_1 \geq q_1^*\left(\bar{\theta}\right), \end{cases} \tag{17}$$

and the bad doctor has no incentive to mimic the good doctor:

$$V(q_1^*(\underline{\theta}), q_2^*(\underline{\theta}) \mid \underline{\theta}) - \alpha(E\left(\theta\right) - \underline{\theta}) \geq \begin{cases} V(q_1, q_2^*\left(\underline{\theta}\right) \mid \underline{\theta}) - \alpha(E\left(\theta\right) - \underline{\theta}) \text{ if } q_1 < q_1^*\left(\bar{\theta}\right) \\ V(q_1, q_2^M\left(\underline{\theta}\right) \mid \underline{\theta}) + \alpha(\bar{\theta} - E\left(\theta\right)) \text{ if } q_1 \geq q_1^*\left(\bar{\theta}\right). \end{cases} \tag{18}$$

The latter can be re-written as: $0 \leq \alpha \leq \tau_1$.

Semi-separating equilibrium. Given the posited strategies, Bayes' Rule can always be applied to $q_1 \in \{q_1^*\left(\bar{\theta}\right), q_1^*\left(\underline{\theta}\right)\}$. Indeed, posterior beliefs when either of these two qualities is observed are $\lambda^S\left(q_1^*\left(\underline{\theta}\right)\right) = 0$ and $\lambda^S\left(q_1^*\left(\bar{\theta}\right)\right) = \frac{\lambda}{1 - r(1-\lambda)}$. These beliefs yield the following expected types: $\theta^S\left(q_1^*\left(\underline{\theta}\right)\right) = \underline{\theta}$, $\theta^S\left(q_1^*\left(\bar{\theta}\right)\right) = \frac{\lambda\bar{\theta} + (1-r)(1-\lambda)\underline{\theta}}{1 - r(1-\lambda)}$. Following similar steps as for the main model we obtain:

$$r^E = 1 - \frac{\lambda}{(1-\lambda)}\left(\frac{\alpha - \tau_1}{\tau_1}\right), \; \lambda^S\left(q_1^*\left(\bar{\theta}\right)\right) = \frac{\tau_1}{\alpha}. \tag{19}$$

Then $\theta^S\left(q_1^*\left(\underline{\theta}\right)\right) = \underline{\theta} + \frac{\tau_1}{\alpha}\left(\bar{\theta} - \underline{\theta}\right)$. The payoff when choosing $q^*\left(\underline{\theta}\right)$, which reveals that the type is low, is given by $V(q^*\left(\underline{\theta}\right) \mid \underline{\theta}) - \alpha(E\left(\theta\right) - \underline{\theta})$. The high type's payoff is $V(q^*(\bar{\theta}) \mid \bar{\theta}) + (\tau - \alpha\lambda)\left(\bar{\theta} - \underline{\theta}\right)$. To summarise, the good doctor always provides high qualities $q_1^*\left(\bar{\theta}\right)$ and $q_2^*\left(\bar{\theta}\right)$. The bad doctor provides qualities $\left\{q_1^*\left(\bar{\theta}\right), q_2^M\left(\underline{\theta}\right)\right\}$ with probability

$r^E = 1 - \frac{\lambda}{(1-\lambda)}\left(\frac{\alpha-\tau_1}{\tau_1}\right)$, and $\{q_1^*(\underline{\theta}), q_2^*(\underline{\theta})\}$ with probability (r^E). The semi-separating equilibrium smoothly connects the separating with the pooling equilibrium.

Appendix 4. Optimal contracting. *Pooling equilibrium.* Formally, the participation constraints (PCs) are $V(q^*(\overline{\theta},p),p,T \mid \overline{\theta}) \geq 0$ for the good doctor and $V(q^*(\overline{\theta},p),p,T \mid \underline{\theta}) \geq 0$ for the bad doctor. The limited liability (L) constraint is the same for both types: $T + pq^*(\overline{\theta},p) - C(q^*(\overline{\theta},p)) \geq 0$. This follows naturally from the assumption that both types have the same contract, the same cost function and provide the same quality under a pooling equilibrium. When the limited liability constraint is binding the payoffs are: $V(q^*(\overline{\theta},p),p,T \mid \underline{\theta}) = \underline{\theta}W(q^*(\overline{\theta},p))$ and $V(q^*(\overline{\theta},p),p,T \mid \overline{\theta}) = \overline{\theta}W(q^*(\overline{\theta},p))$. Substituting for $T = C(q^*(\overline{\theta},p)) - pq^*(\overline{\theta},p)$ into (6), and maximising with respect to the optimal price, we obtain:

$$\left[W_q\left(q^*\left(\overline{\theta},p\right)\right) - C_q\left(q^*\left(\overline{\theta},p\right)\right)\right]\frac{\partial q^*\left(\overline{\theta},p\right)}{\partial p} = 0. \tag{20}$$

The condition suggests that the price should be designed to induce equality between the marginal benefit and marginal cost of quality. Hence $W_q\left(q^*\left(\overline{\theta},p\right)\right) = C_q\left(q^*\left(\overline{\theta},p\right)\right)$, which is satisfied for a unique value for $q^*\left(\overline{\theta},p\right)$, which we refer to as q^o.

Separating equilibrium. The participation constraints are: $V(q^*(\overline{\theta},p),p,T \mid \overline{\theta}) + \alpha(\bar{\theta} - E(\theta)) \geq 0$ for the good doctor and $V(q^*(\underline{\theta},p),p,T \mid \underline{\theta}) - \alpha(E(\theta) - \underline{\theta}) \geq 0$ for the bad doctor; and the limited-liability constraints are: $T + pq^*\left(\overline{\theta},p\right) \geq C(q^*\left(\overline{\theta},p\right))$ and $T + pq^*(\underline{\theta},p) \geq C(q^*(\underline{\theta},p))$. Recall that $\pi(\theta) = T + pq^*(\theta,p) - C(q^*(\theta,p))$ with $\frac{\partial \pi}{\partial \theta} = [p - C_q(q^*(\theta,p))]\frac{\partial q^*(\theta,p)}{\partial \theta} = -\theta B_q(q^*(\theta,p))\frac{\partial q^*(\theta,p)}{\partial \theta} < 0$. The result is analogous to Choné and Ma (2011). The PC of the good and the bad doctor is never binding. This is clearly the case for the good doctor who obtains a positive reputational payoff and zero profits. It is also the case for the bad doctor: the sum of the positive profit and the altruistic component are higher than the negative reputational payoff. If we evaluate the latter at its lowest value, ie at $\alpha = \tau$, we obtain: $V(q^*(\underline{\theta}) \mid \underline{\theta}) - \frac{\left[V(q^*(\underline{\theta})|\underline{\theta}) - V(q^*(\overline{\theta})|\underline{\theta}))\right]}{\overline{\theta}-\underline{\theta}}(E(\theta) - \underline{\theta}) = V(q^*(\underline{\theta}) \mid \underline{\theta})\frac{\overline{\theta}-E(\theta)}{\overline{\theta}-\underline{\theta}} + V(q^*\left(\overline{\theta}\right) \mid \underline{\theta}))\frac{E(\theta)-\underline{\theta}}{\overline{\theta}-\underline{\theta}} > 0$.

Semi-separating equilibrium. Under a semi-separating equilibrium, $\tau(p^S) \leq \alpha \leq \tau(p^P)/\lambda$ we have that the good doctor always provides the higher quality $q^E\left(\overline{\theta},p\right) = q^*\left(\overline{\theta},p\right)$. Instead, the bad doctor provides the low quality $q^*(\underline{\theta},p)$ with probability $r^E = 1 - \frac{\lambda}{(1-\lambda)}\left(\frac{\alpha-\tau(p)}{\tau(p)}\right)$ and the high quality $q^*\left(\overline{\theta},p\right)$ with probability $(1 - r^E)$. The PCs are $V(q^*(\overline{\theta}) \mid \overline{\theta}) + (\tau - \alpha\lambda)\left(\bar{\theta} - \underline{\theta}\right)$ and $V(q^*\left(\overline{\theta}\right) \mid \underline{\theta}) + (\tau - \alpha\lambda)\left(\bar{\theta} - \underline{\theta}\right)$. The L constraints are $T + pq^*\left(\overline{\theta},p\right) \geq C(q^*\left(\overline{\theta},p\right))$ and $T + pq^*(\underline{\theta},p) \geq C(q^*(\underline{\theta},p))$. Again, the PC constraints are never binding. The problem becomes

$$\begin{aligned}\underset{T,p}{Max}\ &\left\{\left(\lambda + (1-\lambda)(1-r^E)\right)\left[W\left(q^*\left(\overline{\theta},p\right)\right) - T - pq^*\left(\overline{\theta},p\right)\right]\right. \\ &\left.+(1-\lambda)r^E\left[W\left(q^*\left(\underline{\theta},p\right)\right) - T - pq^*\left(\underline{\theta},p\right)\right]\right\}\end{aligned} \tag{21}$$

subject to the L constraint for the good doctor $T + pq^*\left(\bar{\theta}, p\right) \geq C(q^*\left(\bar{\theta}, p\right))$. When reputational concerns are sufficiently high we have $r^E = 0$ and the problem reduces to the one solved under a pooling contract. Instead, when reputational concerns are sufficiently low, then $r^E = 1$ and we are back to the separating equilibrium.

Appendix 5. *Proof of Proposition 5.* Let us express the ICT formally first. Beliefs fail the ICT if there exists some q^h such that the following conditions simultaneously hold (recall that in the pooling equilibrium there is neither reputation gain or loss):

$$\underbrace{V\left(q^*\left(\bar{\theta}\right)|\bar{\theta}\right)}_{\text{eq. payoff}} < V\left(q^h|\bar{\theta}\right) + \underbrace{\alpha(\bar{\theta} - E\left(\theta\right))}_{\text{Most fav. beliefs}}, \tag{22}$$

$$\underbrace{V\left(q^*\left(\bar{\theta}\right)|\underline{\theta}\right)}_{\text{eq. payoff}} > V\left(q^h|\underline{\theta}\right) + \underbrace{\alpha(\bar{\theta} - E\left(\theta\right))}_{\text{Most fav. beliefs}}. \tag{23}$$

It is useful to compute τ/λ, that is, the threshold for α above which the pooling equilibrium arises, for the assumed functional forms and price. After substitution in the definition of τ (see equation (3) in Section 2) we obtain $\frac{\tau}{\lambda} = \frac{1}{\lambda}\frac{1}{4}\frac{c^2}{v_2}\frac{\bar{\theta}-\underline{\theta}}{\underline{\theta}}$. We now express conditions (22) and (23) for the assumed functional forms. *Condition (i).* At price p^P and $q = q^*\left(\bar{\theta}\right)$, type $\bar{\theta}$ obtains non-reputational payoff $V\left(q^*\left(\bar{\theta}\right)|\bar{\theta}\right) = T + c\left(1-\bar{\theta}\right)\frac{v_1-c}{2v_2} + \bar{\theta}\left(v_1\frac{v_1-c}{2v_2} - v_2\left(\frac{v_1-c}{2v_2}\right)^2\right) - c\frac{v_1-c}{2v_2}$. This type's non-reputational payoff at any other $q > q^*\left(\bar{\theta}\right)$ becomes $V\left(q|\bar{\theta}\right) = T + c\left(1-\bar{\theta}\right)q + \bar{\theta}v_1 q - \bar{\theta}v_2 q^2 - cq$. By inspection of the expression for $V\left(q^*\left(\bar{\theta}\right)|\bar{\theta}\right)$, the RHS of (22) decreases with q if $q > \frac{v_1-c}{2v_2} = q^*\left(\bar{\theta}\right)$. Therefore condition (i) can be expressed as $q < q^{\max}$, where $q^{\max}$ solves (22) with equality. Once expressions for $V\left(q^*\left(\bar{\theta}\right)|\bar{\theta}\right)$ and for $V\left(q|\bar{\theta}\right)$ are substituted in, the resulting equation has two solutions, but only one of them yields $q > q^*\left(\bar{\theta}\right)$, namely, $q^{\max} = q^*\left(\bar{\theta}\right) + \sqrt{\left(1-\lambda\right)\alpha\frac{\bar{\theta}-\underline{\theta}}{\bar{\theta}v_2}}$. *Condition (ii).* At $p = p^P$ and $q = q^*\left(\bar{\theta}\right)$, type $\underline{\theta}$ obtains non-reputational payoff $V\left(q^*\left(\bar{\theta}\right)|\underline{\theta}\right) = T + c\left(1-\bar{\theta}\right)\frac{v_1-c}{2v_2} + \underline{\theta}\left(v_1\frac{v_1-c}{2v_2} - v_2\left(\frac{v_1-c}{2v_2}\right)^2\right) - c\frac{v_1-c}{2v_2}$. At any other $q > q^o$, type $\underline{\theta}$ obtains non-reputational payoff $V\left(q|\underline{\theta}\right) = T + c\left(1-\bar{\theta}\right)q + \underline{\theta}v_1 q - \underline{\theta}v_2 q^2 - cq$. By inspection of the expression for $V\left(q^*\left(\bar{\theta}\right)|\underline{\theta}\right)$, the RHS of (23) decreases with q if $q > \frac{\underline{\theta}v_1 - c\bar{\theta}}{2\underline{\theta}v_2} \stackrel{def}{=} q_1$. Recall however that we are considering q above $q^*\left(\bar{\theta}\right)$. It is easy to check that $q^*\left(\bar{\theta}\right) > q_1$. Therefore condition (23) can be expressed as $q > q^{\min}$, where $q^{\min}$ solves (23) with equality. Once the expressions for $V\left(q^*\left(\bar{\theta}\right)|\underline{\theta}\right)$ and for $V\left(q|\underline{\theta}\right)$ are substituted in, the resulting equation has again two solutions but only one of them yields $q > q^*\left(\bar{\theta}\right)$. Namely, $q^{\min} = \frac{v_1 - \frac{\bar{\theta}}{\underline{\theta}}c}{2v_2} + \sqrt{\frac{c^2}{4}\left(\frac{\bar{\theta}-\underline{\theta}}{\underline{\theta}v_2}\right)^2 + \left(1-\lambda\right)\alpha\frac{\bar{\theta}-\underline{\theta}}{\underline{\theta}v_2}}$. We now find a necessary and sufficient condition for (i) and (ii) to be compatible. *Conditions (i) and (ii) and the threshold for pooling.* (i) and (ii) are compatible if and only if $q^{\max} > q^{\min}$. It turns out that: (1) $q^{\max} - q^{\min}$ is decreasing in α if and only $\alpha > \frac{1}{4}\frac{c^2}{v_2}\frac{1}{1-\lambda} \stackrel{def}{=} \hat{\alpha}$. (2) $q^{\max} - q^{\min} = 0$

at $\alpha = \frac{\bar{\theta}}{\bar{\theta}-\underline{\theta}} \frac{c^2}{v_2} \frac{1}{1-\lambda} \stackrel{def}{=} \alpha^*$. Notice also that $\alpha^* > \hat{\alpha}$ since $\frac{\bar{\theta}}{\bar{\theta}-\underline{\theta}} > 1 > \frac{1}{4}$. Therefore, (i) and (ii) are incompatible for all $\alpha > \alpha^*$. This implies that OOE beliefs that support the pooling equilibrium pass the ICT test for such α. This proves the proposition.

Proof of Proposition 6. Formally, strategy $\tilde{q}$ is such that

$$\underbrace{V\left(q^*\left(\underline{\theta}\right)|\underline{\theta}\right) - \alpha(E\left(\theta\right) - \underline{\theta})}_{\text{Separating Equilibrium payoff}} = V\left(\tilde{q}|\underline{\theta}\right) + \underbrace{\alpha(\bar{\theta} - E\left(\theta\right))}_{\text{Reputational gain under most favorable beliefs}}. \tag{24}$$

We start by analysing $\tilde{q}$ as defined in the last equation. Total differentiation with respect to α shows that $\tilde{q}$ increases with α. Indeed, $\frac{\partial \tilde{q}}{\partial \alpha} = \frac{-\left(\bar{\theta}-\underline{\theta}\right)}{\underline{V}_q(\tilde{q})}$, where $\underline{V}_q\left(\tilde{q}\right) < 0$ since $\tilde{q} > \underline{q}^*$. Notice that, at $\alpha = \tau$, the expression becomes (after simplification) $\tilde{q}\left(\tau\right) = q^*\left(\bar{\theta}\right)$. (For all $\alpha \leq \tau$, one has the so called 'trivial separation', Cartwright and Patel, 2013). Therefore, for $\alpha > \tau$, the Riley Outcome is given by $q^E\left(\underline{\theta}\right) = \underline{\theta}$ and $q^E\left(\bar{\theta}\right) = \tilde{q} > q^*\left(\bar{\theta}\right)$. In words, as soon as α becomes larger than τ, the high type has to set a higher quality than his non-reputational optimum in order to avoid imitation. We now find the Riley Outcome under the same assumptions in Proposition 5. Solving (24) for $\tilde{q}$ yields

$$\tilde{q} = \frac{1}{2\underline{\theta}v_2}\left(2\sqrt{\alpha\underline{\theta}v_2\left(\bar{\theta}-\underline{\theta}\right)} + \underline{\theta}v_1 - c\bar{\theta}\right).$$

The high type obtains, if he sticks to the Riley Outcome, the best possible beliefs (since the Riley Outcome is separating by definition). However, he is not setting his non-reputational optimum $q^*\left(\bar{\theta}\right)$ if $\alpha > \tau/\lambda$. Therefore, the most favorable out-of-equilibrium belief to sustain the Riley Outcome is that $\theta^s\left(q^*\left(\bar{\theta}\right)\right) = \underline{\theta}$. The next inequality is therefore a necessary condition for the high type not to deviate to $q^*\left(\bar{\theta}\right)$:

$$V\left(q^*\left(\bar{\theta}\right)|\bar{\theta}\right) + \underbrace{\alpha(\underline{\theta} - E\left(\theta\right))}_{\text{Worst possible reput.payoff (negative)}} \leq V\left(\tilde{q}|\underline{\theta}\right) + \underbrace{\alpha(\bar{\theta} - E\left(\theta\right))}_{\text{Reput.gain under the Riley Outcome}}. \tag{25}$$

This can be rewritten as $V\left(q^*\left(\bar{\theta}\right)|\bar{\theta}\right) - V\left(\tilde{q}|\bar{\theta}\right) \leq \alpha\left(\bar{\theta}-\underline{\theta}\right)$. (Incidentally, both sides of last expression increase with α. Indeed, $\tilde{q}'\left(\alpha\right) > 0$ and $V\left(\cdot|\bar{\theta}\right)$ decreases with q when $q > \bar{q}^* < \tilde{q}$. This is why it is impossible to know which effect dominates without further assumptions on V and $\left(\bar{\theta}-\underline{\theta}\right)$). Using the functional forms for W and C given in Proposition 5 and using $p = p^P$, (25) can be rewritten as $\alpha < \alpha^{RO}$ where α^{RO} is given in the proposition.

FIGURES

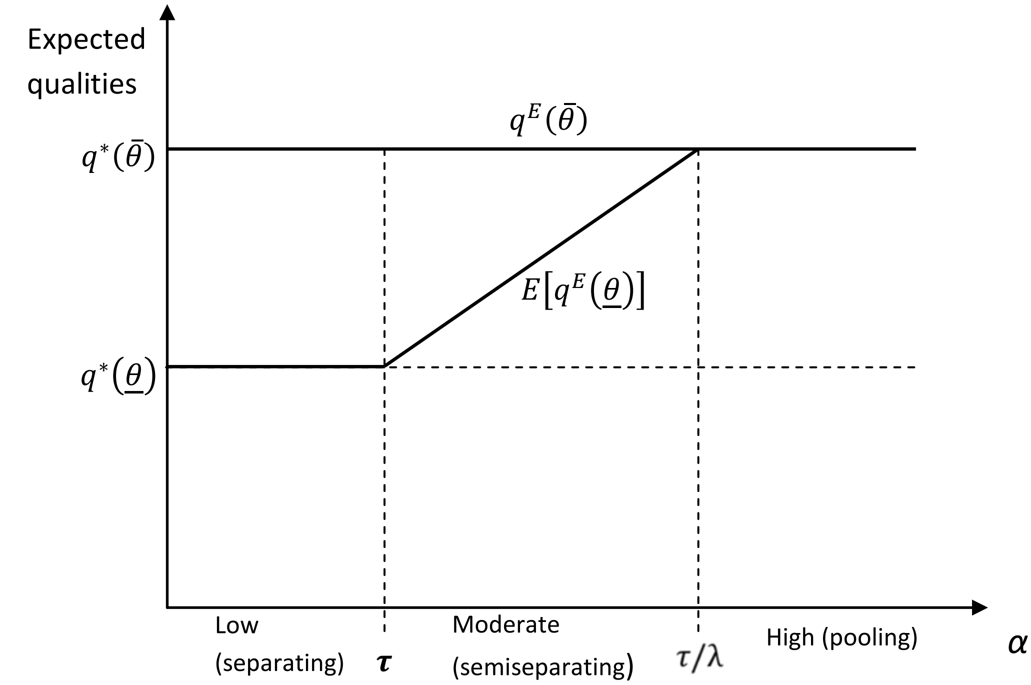

Figure 1. Expected qualities as a function of reputational concerns

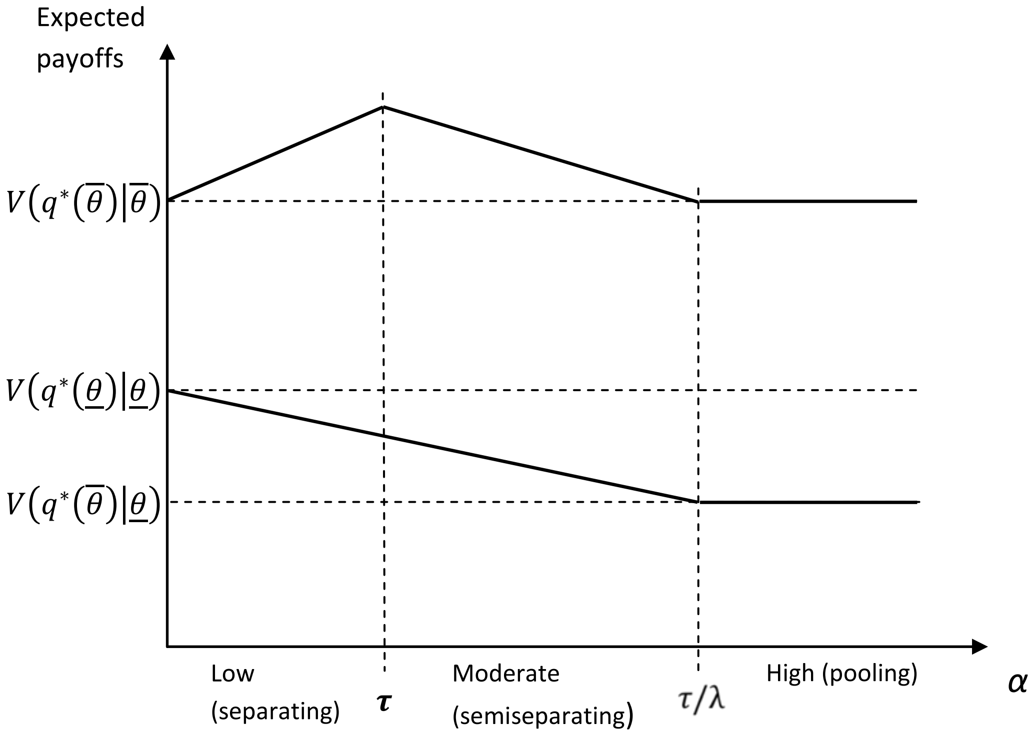

Figure 2. Expected payoffs as a function of reputational concerns

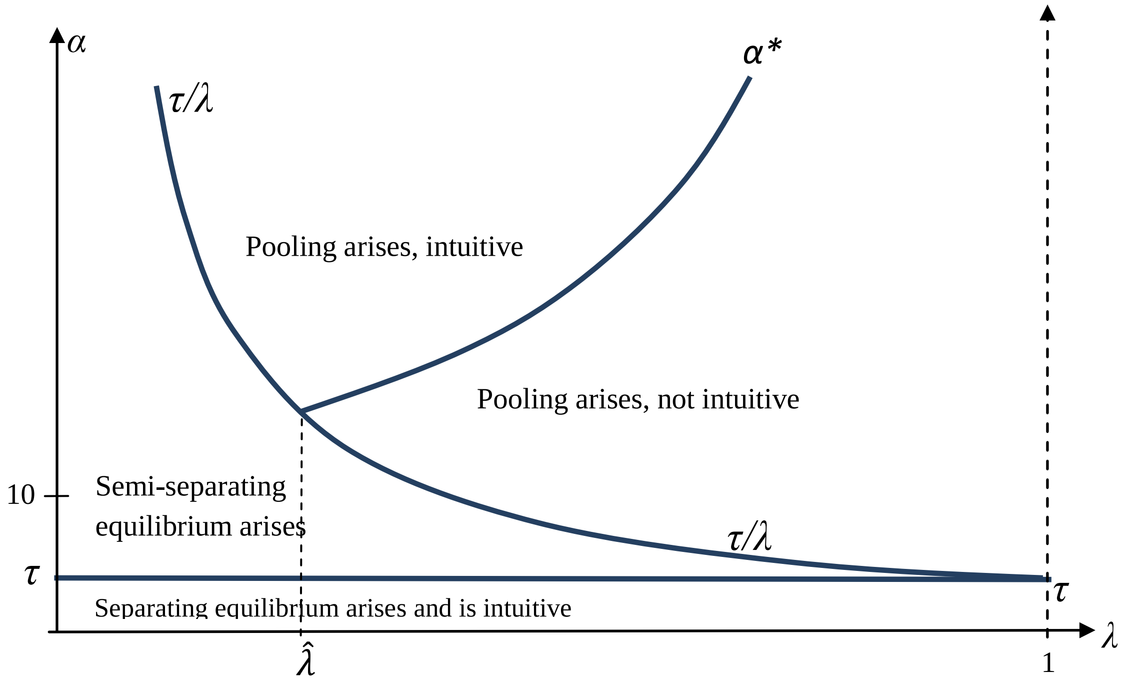

Figure 3. Equilibrium type and the Intuitive Criterion, by proportion of good doctors (λ) and Intensity of reputational concerns (*α)*

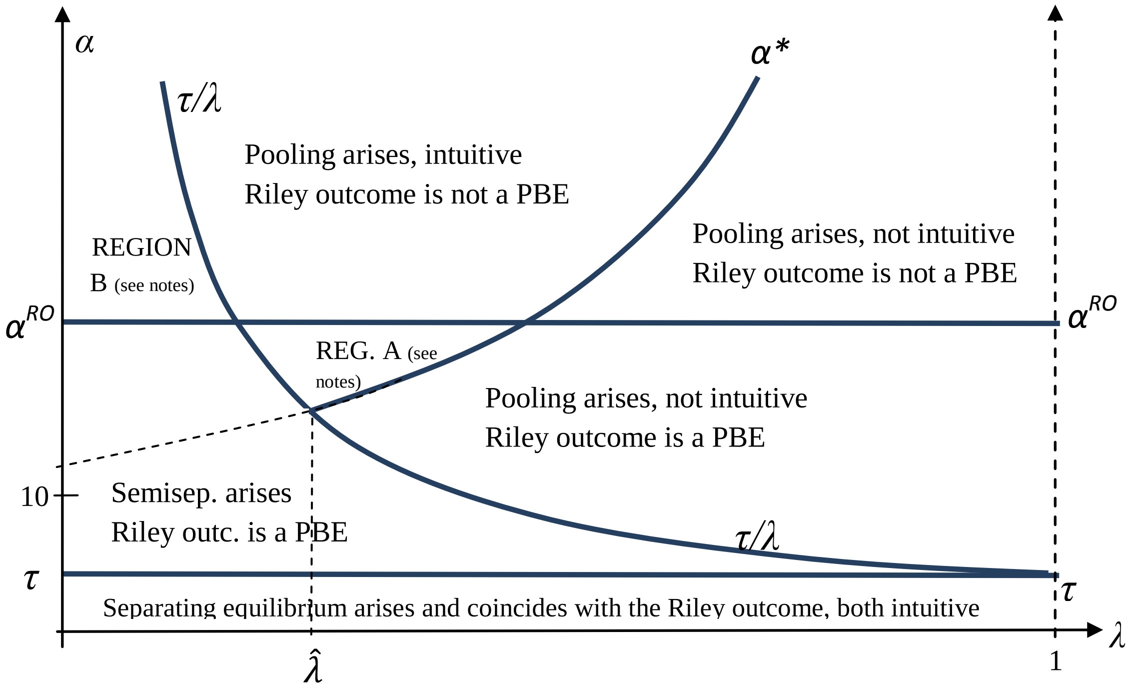

Figure 4. Equilibria and the Riley outcome, by proportion of good doctors (λ) and Intensity of reputational concerns (*α)*

Notes: In Region A, the pooling equilibrium arises and is intuitive. Also the Riley outcome is a PBE. Therefore two PBE equilibria passing the ICT coexist in this region. In Region B, the semiseparating equilibrium arises but the Riley outcome is not a PBE.