Article:

Olivella, Pau and Siciliani, Luigi orcid.org/0000-0003-1739-7289 (2017) Reputational Concerns with Altruistic Providers. Journal of Health Economics. pp. 1-13. ISSN 0167-6296

https://doi.org/10.1016/j.jhealeco.2017.05.003

Reputational Concerns with Altruistic Providers*

Pau Olivella† Luigi Siciliani‡

February 8, 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 intense the less altruistic (*bad*) doctor mimics the more altruistic (*good*) doctor. Otherwise either a separating or a semiseparating equilibrium arises: the bad doctor mimics with probability less than one. Pay-for-performance incentive schemes are unlikely to induce crowding out of observable quality. However, if some dimensions of quality are unobservable, the publication and dissemination of quality indicators will crowd out unobserved quality of the bad doctor. A third-party payer may implement the first-best observable quality by appropriately choosing a single compensation schedule under the pooling equilibrium but not under the separating one.

Keywords: reputation; altruism; doctors, pay for performance; multitasking.

JEL Classifications: I11; I18.

---

*The authors wish to thank 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, 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 (e.g., doctors, hospitals) to provide better care. Providers have at least two sources of motivation: monetary and non-monetary ones. Pay-for-performance incentive schemes can be used to motivate providers and this is the typical economists' focus: 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. The economics literature has long recognised at least two other incentive forces. First, providers may be altruistic: they care about their patients' well-being. Altruism motivates them to provide better quality. Second, providers may care about what other people (their family, the community, their peers, future patients, and so on; henceforth society) think about them. In this study we combine both of these non-monetary forces and seek to explain how they interact with each other. We assume that providers (doctors henceforth) differ in the degree of altruism, which is not directly observable, and they care about their own reputation. For example physicians enjoy being known by society as good doctors, and dislike being known as bad doctors.

The first objective of this study is to investigate the extent to which such reputational concerns can induce bad doctors to provide more quality to avoid a bad reputation. Policymakers increasingly publish, and make available to patients and the general public, information on relative doctors' performance. Examples include the scheme "Quality-Counts" in Wisconsin which compared 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 well-performing doctors and the under-performing ones; they are sometimes (colloquially) known as 'name and shame' where poorly performing doctors are subjected to 'shame' in front of the community. Can the simple fact of publishing information change doctors' behaviour? If so, which doctors change behaviour and in which direction? Do patients and doctors gain from such policies?

We also investigate a number of related 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 when doctors provide different dimensions of quality, some of which are unobservable to society (i.e. within a

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

multi-tasking framework). While some quality dimensions are readily observable (e.g., how much time did the doctor spend with the patient), others may be much harder to observe (e.g. diagnostic effort). Third, we investigate whether there is still scope for publishing and disseminating information on quality even when the third-party payer (a health authority or a health insurer, the purchaser henceforth) can design a pay-for-performance scheme involving a single linear remuneration scheme which is optimally set and maximises expected benefits from quality net of monetary transfers (and subject to doctors' limited-liability constraints).

Our model predicts that policies that publicise doctors' performance, such as 'name and shame', may be virtuous and induce the expected change but only if they imply a sufficiently intense reputational concern. Specifically, if reputational concerns are sufficiently strong, then the bad doctor mimics the good doctor. Formally, a pooling equilibrium arises where all doctors set the good-doctor's non-reputational optimum, where by non-reputational optimum we refer to the quality that a doctor sets in the absence of reputational concerns. However, society cannot update their original beliefs about doctors' true altruism and neither reputational gains nor losses accrue. If reputational concerns are weak, all doctors, good and bad, perform as if no reputational concerns were present. Formally, a separating equilibrium arises where doctors choose their non-reputational optimum, with higher quality provided by the good doctor. The good doctor benefits and the bad doctor suffers from the publication of performance indicators. A semiseparating equilibrium exists for intermediate reputational concerns: the bad doctor randomizes between his non-reputational optimum and that of the good doctor's. This equilibrium continuously connects the outcomes of the separating and pooling equilibria.

Whether the introduction of a pay-for-performance scheme increases or reduces the range of reputational values over which the virtuous pooling equilibrium arises, is in principle indeterminate. One the one hand, higher prices increases the good doctor's performance and makes it more costly for the bad doctor to mimic, which in turn tends to favour crowding out. On the other hand, raising prices increases overall revenues more when performance is high and makes it more attractive to the bad doctor to mimic, which instead 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 than 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 therefore 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, we do obtain another source of crowding out. If some dimensions of quality cannot be observed by the patient and society (i.e. in the presence of multitasking), then publicizing the quality performance in the observable dimension can induce a pooling equilibrium that is detrimental for the patient. When the bad doctor mimics the observable quality he finds advantageous to reduce the non-observable quality to a level which is below the one that would maximize his non-reputational payoff. Name and shame policies may therefore *crowd out* non-observable dimensions of quality. If these dimensions are important, patient's benefit may overall reduce as a result of such policies.

As for the design of optimal pay-for-performance schemes, we show that only if reputational concerns are strong a single (and linear) remuneration scheme suffices to induce all doctors to set the allocative efficient quality, that is, the quality that maximizes benefits minus costs. Allocative efficiency for both doctors' type is instead not obtained for weak reputational concerns. Intuitively, under strong reputational concerns the payer implements allocative efficiency for the good doctor, and by pooling accomplishes allocative efficiency of the bad doctor as well. In contrast, under weak reputational concerns a separating equilibrium arises and it is not possible to achieve allocative efficiency for both doctors simultaneously. In summary, if the purchaser is constrained by the use of a single linear scheme, policies which publicize quality can make patients and purchasers better off despite the payment being optimally set. This result is relevant for policy and suggests that pay-for-performance policies are not a substitute for policies aimed at disseminating quality indicators.

There have been some studies aimed at quantifying the effects of publicizing performance reports. Some of these provide evidence of pooling. For instance, 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, while it is true that "public report" hospitals did not immediately pool at high quality, these hospitals were more likely to do so than those hospitals whose performance was not publicized. Similarly, Fichera *et al.* (2014), in their survey on relative performance evaluation experiments, report 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." The same authors point out that "[p]roviders were often given data on their own performance and the performance of the average provider or their rank in the distribution of performance over anonymized providers. Later, these data were deanonymized and sometimes publicly reported." (p. 113) Wang et al. (2011) examine the impact of coronary bypass report cards. The find that poor performing hospitals or surgeons responded with a reduction in volume, while highly rated hospitals and surgeons did not respond.

1.1 Related Literature

The literature on altruism and intrinsic motivation is now vast, both in the empirical and the theoretical front. Within the public and health economics literature the assumption of motivated agents is shared by many studies.[2] Establishing that reputational concerns matter has also been investigated. 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 recent review).

However, few studies formally include the possibility that reputational concerns come from society learning about doctor's altruism from observed actions. These studies can be classified into two groups. In the first, this effect is either directly assumed in the doctor's payoff function (Siciliani, 2008) or comes about from the implicit assumption that individuals with different altruism (types) choose different actions, and altruism has a continuous support (Bénabou and Tirole, 2006). In the second, reputational concerns are explicitly modelled as a formal signalling game, as we do here. Bénabou and Tirole (2011) exploit a model which, differently from ours, has a continuous type space and the signal is dichotomous. Moreover, their focus is on analyzing and comparing sources of crowding out that are very different from ours (which we outline below). Much closer to us are the works by Jeitschko and Normann (2012) and Cartwright and Patel (2013). These authors consider, as we do here, signals with continuous support and dichotomous types. Jeitschko and Normann (2012), an experimental game theory exercise, provide an extensive characterization of the perfect Bayesian equilibria of a model that, although close to ours, is much more restrictive. Cartwright and Patel (2013) present a model of fundraising and donations. They however concentrate on equilibria that are very different from the ones we focus on. Indeed, 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). This outcome consists on the separating equilibrium that is least costly to the (in our terminology) good doctor.[3] We instead propose a simpler

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

[3] Moreover, the focus of their paper is different: they are interested in checking whether reporting

equilibrium selection procedure, based on more naive beliefs. Basically, we first posit that doctors choose to perform as if no reputational concerns were present and then check whether, when taking into account reputational effects, they want to deviate from the posited choice. The equilibria outlined (separating, semiseparating, and pooling) can then be supported by out-of-equilibrium beliefs satisfying monotonicity (observing better-than-equilibrium performance leads to no worse beliefs) and pessimism (worst possible belief for the doctor subject to the previous monotonicity condition). We then 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. Briefly, both implications are true if (i) the function translating doctor's care ("care" being the observable signal) into patients benefits is quadratic and (ii) reputational concerns are sufficiently intense. Condition (i) is particularly justified in the setup we are analyzing. For instance, consider the case where doctor's care is measured by the number of diagnostic tests he performs on the patient (e.g., X-rays), which becomes harmful if excessive.[4] In the donations literature, in contrast, it is more natural to assume that a larger donation brings a higher direct benefit to its beneficiary. Chen (2011) also investigates a signalling model where doctors with low quality may select low-severity patients to disguise as high-quality doctors. He shows that if doctors face the same distribution of patients' type, low-quality doctors have no incentive to select patients, while this arises when they face different distributions of patients' type. Rodriguez-Barraquer and Xu (2015) have agents seek promotion by choosing a difficult task. These authors also obtain a pooling equilibrium under some conditions.

As mentioned above, the possibility that pecuniary incentives crowd out reputational concerns only occurs, within the contest of our model, under very stringent conditions. However, such 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 Ariely *et al.*, 2009). In psychology (Deci, 1975) and in particular in Self Determination Theory, Ryan and Deci (2000) argue: "[r]esearch revealed that not only tangible rewards but also threats, deadlines, directives, pressured evaluations, and imposed goals diminish intrinsic motivation because, like tangible rewards, they conduce toward an external perceived locus of causality."(p. 70). These theories, and related ones, have more to do with (suppressing) the direct pleasure that one's action produces rather than with concerns about one's reputation of being altruistic. In cases where the true objectives of

actions to society in categories (say high, medium or low, "category reporting") rather than reporting actions directly ("exact reporting") changes the equilibrium actions.

[4] Real-world examples of measures of doctor's performance that could potentially lead to counterproductive effects can be found in the list of quality measures used in the AQ Program in North West England in 2008 (Fichera *et al.*, 2014). For instance, for patients with acute myocardial infarction, the list includes: Aspirin at arrival; aspirin at discharge; beta blocker at arrival; and beta blocker at discharge.

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, Falk and Kosfeld, 2006; Fehr and List, 2004; Fehr and Falk, 2002; 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 patient's well-being. However, we do obtain another source of crowding out. If some dimensions of quality cannot be observed by the patient and society, i.e. in the presence of multitasking, then publicizing the quality performance in the observable dimension can induce the bad doctor to reduce and therefore *crowd out* non-observable dimensions of quality. This form of crowding has been obtained in previous literature (Eggleston, 2005; Kaarboe and Siciliani, 2011) as a result of the introduction of pay-for-performance schemes, but not as a result of name and shame policies.

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 where one dimension of quality cannot be observed and provide a novel source of crowding out. In Section 5 we characterize the optimal (linear) remuneration contract. In Section 6 we provide conditions for our equilibria to pass the Intuitive Criterion test and for the Riley outcome not being an equilibrium. Section 7 concludes. Most technical derivations are relegated to the Appendix.

2 The model

The players are a doctor, a third-party payer (e.g. a public or private insurer), and society (e.g. patients, their family and friends, provider'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 $\theta \in \{\underline{\theta}; \overline{\theta}\}$ with $\underline{\theta} < \overline{\theta}$. We refer to the more altruistic doctor (with $\theta = \overline{\theta}$) as the *good doctor*. We refer to the less altruistic provider (with $\theta = \underline{\theta}$) as the *bad doctor*. The prior probability that the doctor is good is common knowledge and equal to $\lambda \geq 0$. We assume that altruism is private information. Both types have the same costs of delivering quality, given by $C(q)$, with $C_q > 0$ and $C_{qq} \geq 0$. These costs are the sum of the monetary and/or non-monetary ones (e.g., diagnostic effort, opportunity costs of time spent with the patient, and so on).

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, for example, in Ellis and McGuire, 1986). Formally, patients derive benefits $W(q)$, with $W_q(q) \geq 0$ if $q \leq \widehat{q}$, possibly with $\widehat{q} = \infty$, $W_q(q) \leq 0$ if $q \geq \widehat{q}$, and $W_{qq} \leq 0$. The marginal benefit may become negative if unnecessary tests and X-rays are prescribed, or drugs with side effects and no health gains. For brevity we refer to q as quality in the rest of the paper.

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

Doctor's preferences are represented by a linear and additively separable utility function over money, an altruistic component, and reputational concerns. The revenues are given by $T + pq$, where T is a lump-sum payment and p is bonus for additional quality (e.g. as part of a 'pay-for-performance' scheme).[5] His (monetary) profits are $\pi(q) = T + pq - C(q)$.

Similarly to Ellis and McGuire (1986) and Chalkley and Malcomson (1998), altruism is expressed as fraction of patients' benefits, $\theta W(q)$. Hence, θ corresponds to "pure altruism" (Andreoni, 1989), in the sense that the doctor cares about patients' wellbeing. In the special case where $W(q) \equiv q$, parameter θ can be interpreted as the degree of intrinsic motivation of the provider, as in Dixit (2005) and Besley and Ghatak (2005). The sum of the first two components is defined with

$$V = \pi + \theta W, \tag{1}$$

which we refer to as the *non-reputational payoff.*

The reputational concerns convey that the doctor cares about society's *impression* of his own altruism. This impression comes from the composition of two elements. The first element measures how intensely either society, or the doctor himself, takes into account this impression. To be precise, let parameter α_0 be determined by either the doctor's own preferences (how much he cares about what others think of him) or by society's preferences (how much society cares about altruism). Then let α_1 be the number of people who directly

[5] Examples include the Medicare Programme in the United States, which rewards financially 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).

learn, without any policy intervention, about the quality of his care (e.g. through family, friends, word of mouth, and social networks) and α_2 captures the amplification effect of publishing and dissemination of scorecards or other quality indicators. We summarise these three forces in a single parameter $\alpha \equiv \alpha_0 (\alpha_1 + \alpha_2)$. For instance, if the doctor does not care at all about what others think of him and society does not care about doctor's altruism then $\alpha_0 = 0$ and neither α_1 nor α_2 matter. If the doctor does care about others' perceptions (be as a result of his own character or as a reduced-form of some future payoff) then, even in the absence of any public intervention ($\alpha_2 = 0$), he may still have some reputational concerns, since $\alpha = \alpha_0 \alpha_1 > 0$. The second element of the reputational concerns is how society's impression is determined. We assume that it is given by the difference between the conditional expectation θ^S, based on the posterior beliefs $\lambda^S(q)$, and the unconditional average $E(\theta)$. This is consistent with the idea that if no new and relevant information is revealed, the provider's reputation remains the same. For instance, in a pooling equilibrium where all doctors provide the same quality, observed quality is not informative and there should be no reputational gain or loss. Formally, reputational concerns generate the payoff $\alpha(\theta^s(q) - E(\theta))$. If observing some (low) quality generates posterior beliefs that the doctor is bad (if $\lambda^s(q) = 0$) the reputational payoff is negative. If observing quality is uninformative then $\lambda^s(q) = \lambda$ and there is no reputation gain or loss.

To sum up, doctors' preferences are represented by

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

It is useful to define $q^*(\theta)$ as the optimal quality of doctor θ facing no reputational concerns (when $\alpha = 0$). Hence $q^*(\theta)$ maximises $V(q \mid \theta)$ and satisfies $p + \theta W_q = C_q$. We assume the Second Order Condition is satisfied: $\theta W_{qq} - C_{qq} < 0$. The marginal benefit from quality due to monetary and altruistic concerns is equal to the marginal cost. We refer to $q^*(\theta)$ as the *non-reputational optimum quality for type* θ.

Our equilibrium concept to solve the model is the Perfect Bayesian Equilibrium (PBE).

Definition 1 *An equilibrium is a pair of functions of qualities* $q^E(\theta) : \{\underline{\theta}; \overline{\theta}\} \to R_+$ *and beliefs* $\lambda^S(q) : R_+ \to [0,1]$ *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* $\theta^S(q) = \lambda^S(q)\bar{\theta} + (1 - \lambda^S(q))\underline{\theta}$,
(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 (i.e., 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. 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 our beliefs to satisfy the following properties:

(1) [Monotonicity] The beliefs assigned to equilibrium actions –determined by Bayes' rule–, together with the beliefs assigned to out of equilibrium actions, must define a function that is not decreasing in q;

(2) [Pessimism] Beliefs assigned to out of equilibrium actions are the most pessimistic subject to (1).

Unfortunately, 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, an issue to which we return to in Section 7. Rather than applying this criterion directly, we focus on equilibria where the good doctor takes his non-reputational optimum. This gives rise to pooling equilibria where also the bad doctor sets the good doctor's reputational optimum. In Section 7 we provide conditions under which this equilibrium passes the Intuitive Criterion test.

In the next three sections we characterise the set of parameter values which sustain a separating, a pooling and a semi-separating equilibrium, respectively.

2.1 Separating equilibrium

Each provider's type θ chooses quality $q^*(\theta)$, which maximises the non-reputational payoff. The good doctor provides the higher quality, $q^*\left(\bar{\theta}\right) > q^*\left(\underline{\theta}\right)$. Then beliefs in the equilibrium path are such that observing a high (low) quality signals with certainty high (low) altruism: $\lambda^S\left(q^*\left(\bar{\theta}\right)\right) = 1$ and $\lambda^S\left(q^*\left(\underline{\theta}\right)\right) = 0$. It can be shown that the good doctor never has an incentive to mimic the bad doctor. Instead, the bad doctor has no incentive to mimic the good doctor if the reputational gains are sufficiently small (see Appendix 1):

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

The parameter τ has an intuitive interpretation: it conveys the cost of the bad doctor of disguising as a good doctor, relative to the distance between the two types.

In terms of payoffs, the good doctor enjoys

$$V(q^*(\bar{\theta}) \mid \bar{\theta}) + \alpha(\bar{\theta} - E(\theta)), \tag{4}$$

which includes an increase in reputation, while the bad doctor enjoys

$$V(q^*(\underline{\theta}) \mid \underline{\theta}) - \alpha(E(\theta) - \underline{\theta}), \tag{5}$$

which includes a loss of reputation.

We summarise with the following proposition.

Proposition 2 *If reputational concerns are sufficiently small, a separating equilibrium arises where the good doctor provides a higher quality than the bad doctor. Formally, for* $0 \leq \alpha \leq \tau$*, in equilibrium we have* $q^E\left(\overline{\theta}\right) = q^*\left(\overline{\theta}\right)$ *and* $q^E\left(\underline{\theta}\right) = q^*\left(\underline{\theta}\right)$*. The good doctor enjoys a reputation gain while the bad doctor suffers a reputation loss.*

2.2 Pooling equilibrium

Suppose that both types of doctor choose the high quality: $q^E\left(\theta\right) = q^*(\bar{\theta})$ 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 after observing the high quality $q^*\left(\bar{\theta}\right)$. Hence $\lambda^S\left(q^*\left(\bar{\theta}\right)\right) = \lambda$ and the expected type conditional on patients observing the high quality is the average type, $\theta^s\left(q^*\left(\bar{\theta}\right)\right) = E\left(\theta\right)$. It can be shown that 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 sufficiently large (see Appendix 1):

$$\alpha \geq \frac{\tau}{\lambda}. \tag{6}$$

The 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. No type enjoys a reputation gain or loss. This is unsurprising. Observing the high quality $q = q^*(\overline{\theta})$ is not informative since both types are choosing the same quality. Although the good doctor is setting its optimal quality, the bad doctor suffers a mimicking cost of $V(q^*\left(\underline{\theta}\right) \mid \underline{\theta}) - V(q^*(\overline{\theta}) \mid \underline{\theta})$. He is willing to incur this cost to avoid a reputation loss.

We summarise with the following proposition.

Proposition 3 *If reputational concerns are sufficiently high, a pooling equilibrium arises where both doctors choose a high quality and neither type gains or loses any reputation. Formally, for* $\alpha \geq \tau/\lambda$*, we have* $q^E\left(\overline{\theta}\right) = q^E\left(\underline{\theta}\right) = q^*\left(\overline{\theta}\right)$*. Doctors do not enjoy a reputational gain neither a loss.*

Comparing propositions 2 and 3, there is an empty intersection for intermediate levels of reputational concerns, namely $\tau < \alpha < \tau/\lambda$, 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 potential reputation gain (given by $\alpha(\overline{\theta} - E\left(\theta\right))$. He must content himself with avoiding a reputation loss ($\alpha\left(E\left(\theta\right) - \underline{\theta}\right)$). Under pooling individuals cannot infer the level of altruism from the observed qualities.

2.3 Semi-separating equilibrium

We now derive the equilibrium for $\tau < \alpha < \tau/\lambda$, which turns out to be semi-separating. Suppose that the good doctor chooses the high quality $q^E(\bar{\theta}) = q^*(\bar{\theta})$ with certainty, and the bad doctor chooses the low quality with probability r and the high quality with probability $(1-r)$:

$$q^E(\underline{\theta}) = \begin{cases} q^*(\bar{\theta}) \text{ with probability } 1-r, \\ q^*(\underline{\theta}) \text{ with probability } r. \end{cases} \tag{7}$$

Then the equilibrium is characterised by (see Appendix 1):

$$r^E \stackrel{\text{def}}{=} 1 - \frac{\lambda}{(1-\lambda)}\left(\frac{\alpha-\tau}{\tau}\right) < 1 \tag{8}$$

and $\lambda^s(q^*(\bar{\theta})) = \frac{\tau}{\alpha}$; $\lambda^s(q^*(\underline{\theta})) = 0$.

The equilibrium payoffs for the good and bad doctor are, respectively, $V(q^*(\bar{\theta}) \mid \underline{\theta}) + (\tau - \alpha\lambda)(\bar{\theta} - \underline{\theta})$ and $V(q^*(\bar{\theta}) \mid \bar{\theta}) + (\tau - \alpha\lambda)(\bar{\theta} - \underline{\theta})$. In equilibrium, the probability of facing a good doctor conditional on observing a low quality is zero. The probability of facing a good doctor conditional on observing a high quality is positive but less than one. Hence the observation of $q^*(\underline{\theta})$ induces the sure belief that the provider is a bad doctor and the consequent reputation loss. In contrast, the observation of $q^*(\bar{\theta})$ is not fully informative. It could either come from a good doctor or be the outcome of the randomization performed by the bad doctor, who would then be mimicking the good doctor. The posterior probability of observing the good doctor upon the observation of $q^*(\bar{\theta})$ is larger than the expected type, and the bad doctor enjoys a reputation gain, although small. This reputation gain is given by $(\tau - \alpha\lambda)(\bar{\theta} - \underline{\theta})$, which is decreasing in reputational concerns α. This is not surprising since the bad doctor tends to mimic with a higher probability when reputational concerns are larger. The payoff of the good doctor also contains the small reputation gain $(\tau - \alpha\lambda)(\bar{\theta} - \underline{\theta})$.

We summarise with the following proposition.

Proposition 4 *If reputational concerns are intermediate, a semi-separating equilibrium arises where the good doctor provides his non-reputational optimum and the bad doctor randomizes between his non-reputational optimum and disguising. The expected difference in qualities between the two types is smaller than under the separating equilibrium and decreases with the intensity of reputational concerns. Formally, for $\tau < \alpha < \tau/\lambda$, we have $q^E(\bar{\theta}) = q^*(\bar{\theta})$ and $q^*(\underline{\theta}) < E(q^E(\underline{\theta})) < q^*(\bar{\theta})$. Payoffs decrease with the intensity of reputational concerns.*

It can be shown that the semi-separating smoothly connects the separating equilibrium with the pooling equilibrium (see Appendix 1). This is illustrated in Figure 1 (qualities) and Figure 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 bad reputation. When reputational concerns are strong, the desire to avoid the bad reputation is so important that the bad doctor provides the same quality as the good doctor.

The comparison of payoffs across different equilibria provide some interesting insights. The bad doctor is always better off under a separating than under a pooling equilibrium. When reputational concerns are weak, 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 strong enough to induce a pooling equilibrium, patients and society cannot distinguish at all between the 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 to just avoid a bad reputation, instead of gaining a good one.

Figures 1 and 2 also show how an increase in the proportion of good doctors λ increases the range over which the pooling equilibrium arises. The presence of more good doctors increases the expected altruism $E(\theta)$ under the pooling equilibrium and the reputational gain. In turn, this makes the bad doctor more willing to mimic the good doctor. A larger proportion of good doctors also implies that under the semi-separating equilibrium the bad doctor chooses the high quality with a higher probability, and reduces the range of reputational concerns over which a semi-separating equilibrium arises.

[Figures 1 and 2 here]

3 Crowding in or crowding out?

As explained in the introduction, one form of 'crowding out' appears if an increase in pecuniary remuneration weakens the effect of the reputational concerns. Since α is exogenous, in our model this can only occur if variations in prices p, due for example to the introduction of a 'pay-for-performance' scheme, affect the equilibrium. Recall that increases in τ implies a smaller range of reputational concerns over which the pooling equilibrium arises and that τ gives the cost of the bad doctor of disguising as the good doctor relative to the distance between the two types.

Does an increase in price generate crowding out or crowding in (i.e. increases or reduces τ)? We obtain (see Appendix 2):

$$\frac{\partial \tau}{\partial p} = \underbrace{-\frac{q^*(\overline{\theta}) - q^*(\underline{\theta})}{\overline{\theta} - \underline{\theta}}}_{\text{revenue effect}} + \underbrace{W_q(q^*(\overline{\theta}))\frac{\partial q^*(\overline{\theta})}{\partial p}}_{\text{quality effect}} \tag{9}$$

The effect is in general indeterminate. Intuitively, higher prices increase the revenues when the bad doctor provides the lower quality $q^*(\underline{\theta})$ and when he provides the high quality $q^*(\overline{\theta})$. However, since revenues are higher when the high quality is provided, a higher price tends to increase the utility of the bad doctor more when he is disguising as the good doctor compared to when not (first term of the above expression). This effect, which we call *revenue* effect, is negative and tends to reduce τ. However, a higher price also increases both qualities $q^*(\underline{\theta})$ and $q^*(\overline{\theta})$. By the envelope theorem we know that an increase in $q^*(\underline{\theta})$ will have no effect on the utility of the bad doctor, while an increase in $q^*(\overline{\theta})$ will reduce it because it brings the bad doctor even further away from his non-reputational optimum quality. This effect, which we call *quality* effect, is positive and tends to increase τ.

Which of the two effects dominates ultimately depends on the sign of $\partial^2 q^*(\theta)/\partial^2\theta$, which can simply be interpreted as whether the good doctor proportionally provides higher or lower quality than the bad doctor. In Appendix 2, we prove the following proposition.

Proposition 5 *Crowding in (out) arises if the good doctor proportionally provides lower (higher) quality than the bad doctor, i.e. 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 (formally, a sufficient, but not necessary condition is that $W_{qq} < 0$, $W_{qqq} \leq 0$ and $C_{qqq} \geq 0$). This is always the case if the cost function is quadratic or linear (so that $C_{qqq} = 0$). For crowding out to arise, the marginal benefit from quality has to be constant or mildly decreasing and the convexity of the cost function has to sufficiently increase with quality (for example cost is exponential).

The key policy insight is that under a wide range of scenarios (e.g. decreasing marginal benefit) crowding out seems unlikely to arise. Therefore, policies which introduce pay-for-performance scheme do not seem to be in conflict with the introduction of report cards. This conclusion is however valid when report cards and pay-for-performance schemes capture all dimensions of patients' 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. Define q_1 as quality dimension 1 which can be observed by the patient and society and q_2 as quality dimension 2 which cannot be observed by society. The patient lacks the expertise to evaluate q_2 or is unable to report to society. Potential reputational payoffs for the doctors might therefore arise by acting on quality 1, while changes in quality 2 are inconsequential for doctor's reputation.

Like in the main model, qualities are chosen by the doctor, and altruism θ takes two possible values $\{\underline{\theta}; \overline{\theta}\}$ with $\underline{\theta} < \overline{\theta}$. Both doctor types have the same costs of quality, given by $C(q_1, q_2)$, with $C_{q_1} > 0$, $C_{q_2} > 0$, $C_{q_1 q_1} > 0$ and $C_{q_2 q_2} > 0$. Critically, we assume that the two quality dimensions are *substitutes,* $C_{q_1 q_2} > 0$ with $C_{q_1 q_1} C_{q_2 q_2} > C^2_{q_1 q_2}$ (the latter ensures the problem is well behaved). An increase in quality in one dimension increases the marginal cost of quality in the other dimension. This may lead to what is commonly known as the *multitasking* problem: incentivizing one dimension of quality may trigger a reduction in the unincentivised dimension of quality (Holmstrom and Milgrom, 1991; Eggleston, 2008; Kaarboe and Siciliani, 2011). In this section we identify the scenarios under which the multitasking problem arises.

We assume that patients derive benefits from both dimensions of quality given by $W(q_1, q_2)$, with $W_{q_1} > 0$, $W_{q_2} > 0$ and $W_{q_1 q_2} = 0$. Benefits are therefore separable in quality. Making qualities complements or substitutes in benefit would make the model more complex without altering the key insights (see Kaarboe and Siciliani, 2011, for a model without reputational concerns). Doctor's revenues are given by $T + pq_1$, where p is a unit price for observable quality 1. Since quality 2 is not observable, there is no pay-for-performance scheme which can be designed in relation to quality 2.

Profits are $\pi(q_1, q_2) = T + pq_1 - C(q_1, q_2)$. Altruism is a fraction of patients' benefits, $\theta W(q_1, q_2)$. The non-reputational payoff is $V = \pi + \theta W$. Reputational concerns arise only as a result of changes in the observable quality 1 generating a payoff equal to $G(q_1) = \alpha(\theta^s(q_1) - E(\theta))$ and do not depend on the unobservable quality 2, where $\theta^s(q_1)$ is society's posterior belief once q_1 has been observed.

Define the functions $q_1^*(\theta)$ and $q_2^*(\theta)$ as the *non-reputational* optimal qualities of doctor θ facing no reputational concerns (when $\alpha = 0$). These are derived by maximising $V(q_1, q_2 \mid \theta)$:

$$p + \theta W_{q_1}(q_1^*) = C_{q_1}(q_1^*, q_2^*), \qquad (10)$$

$$\theta W_{q_2}(q_2^*) = C_{q_2}(q_1^*, q_2^*). \qquad (11)$$

The marginal benefit of raising quality due to monetary and altruistic concerns is equal to the marginal cost. The problem is well behaved (Appendix 3). An increase in price, or the introduction of a pay-for-performance scheme, incentivizes the observable quality but disincentivizes unobservable quality, an illustration of the multitasking problem, $\partial q_1^*/\partial p > 0, \partial q_2^*/\partial p < 0$. If the degree of cost substitution in qualities is not too high, then both quality dimensions are increasing in altruism, $\partial q_1^*/\partial \theta > 0, \partial q_2^*/\partial \theta > 0$ (Appendix 3). As for the main model, our equilibrium concept is the PBE (Appendix 3).

Separating equilibrium. Each doctor's type θ chooses qualities $q_1^E(\theta) = q_1^*(\theta)$ and $q_2^E(\theta) = q_2^*(\theta)$ that maximise non-reputational payoffs.[6] For these strategies and beliefs to constitute a separating Bayesian Equilibrium, we need the incentive-compatibility constraints for both types to be satisfied. The good doctor never has an incentive to mimic the bad doctor. Instead, the bad doctor has no incentive to mimic the good doctor only if the reputational gains captured by mimicking are sufficiently low, with

$$0 \leq \alpha \leq \tau_1 \stackrel{\text{def}}{=} \frac{V(q_1^*(\underline{\theta}), q_2^*(\underline{\theta}) \mid \underline{\theta}) - V(q_1^*\left(\overline{\theta}\right), q_2^*\left(\overline{\theta}\right) \mid \underline{\theta}))}{\overline{\theta} - \underline{\theta}}. \tag{12}$$

See Appendix 3 for details. We summarise with the following proposition.

Proposition 6 *If reputational concerns are sufficiently small, a separating equilibrium arises where both types of doctor chose qualities which maximise the non-reputational payoff. Formally, for $0 \leq \alpha \leq \tau_1$, in equilibrium we have $q_1^E\left(\overline{\theta}\right) = q_1^*\left(\overline{\theta}\right)$, $q_2^E\left(\overline{\theta}\right) = q_2^*\left(\overline{\theta}\right)$ and $q_1^E(\underline{\theta}) = q_1^*(\underline{\theta})$, $q_2^E(\underline{\theta}) = q_2^*(\underline{\theta})$.*

The results under a separating equilibrium are therefore qualitatively similar to those obtained in Section 2. This is not the case for the pooling equilibrium as shown below.

Pooling equilibrium. The good doctor chooses both the observable and unobservable qualities which maximise the non-reputational payoff: $q_1^E\left(\overline{\theta}\right) = q_1^*\left(\overline{\theta}\right)$, $q_2^E\left(\overline{\theta}\right) = q_2^*\left(\overline{\theta}\right)$. The bad doctor, driven by the reputational concerns, chooses the same observable quality 1 provided by the good doctor: $q_1^E(\underline{\theta}) = q_1^*(\overline{\theta})$ (Appendix 3). Since qualities are substitutes, the bad doctor compensates by providing a low level of quality in the unobserved dimension. Formally, the bad doctor's choice of quality 2 is such that:

$$q_2^M(\underline{\theta}): \quad \underline{\theta} W_{q_2}\left(q_2^M\right) = C_{q_2}(q_1^*(\overline{\theta}), q_2^M), \tag{13}$$

which implies $q_2^E(\underline{\theta}) = q_2^M(\underline{\theta}) < q_2^*(\underline{\theta})$:[7] the unobservable quality chosen by the bad doctor under a *pooling* equilibrium is lower than the quality chosen by the bad doctor

[6] Beliefs in the equilibrium path are such that observing a high (low) quality signals with certainty high (low) altruism: $\lambda^S\left(q_1^*\left(\overline{\theta}\right)\right) = 1$ and $\lambda^S\left(q_1^*(\underline{\theta})\right) = 0$, where recall that only q_1 is observable. Out of equilibrium beliefs should satisfy monotonicity and pessimism, so $\lambda^S(q_1) = 0$ (so $\theta^s(q_1) = \underline{\theta}$) for any $q < q_1^*(\underline{\theta})$ and $\lambda^S(q_1) = 1$ (so $\theta^s(q_1) = \overline{\theta}$) for any $q_1 \geq q_1^*(\underline{\theta})$.

[7] The proof is straightforward. In the choice of quality 2 for the bad doctor, the marginal benefit of quality is the same under the separating and the pooling equilibrium, but the marginal cost of quality is

under a *separating* equilibrium. This is the key result of this section. In a multitasking framework, the bad doctor provides the same observable quality of the good doctor but then compensates by providing lower quality in the unobservable dimension. Therefore, differently from the main model, it is not necessarily the case that patients can benefit from policies which enhance reputation by publishing performance indicators on quality (in our model a large increase in α which switches the equilibrium from separating to pooling). The increase in patients' benefit arising from the higher observable quality could be offset by the reduction in patients' benefit from a reduction in the unobservable quality.

In terms of incentive-compatibility constraints, it is still the case that the good doctor has no incentive to mimic the bad doctor. The bad doctor under a pooling equilibrium does have an incentive to mimic the good doctor when the reputational gain is sufficiently high and the following condition is satisfied:

$$\alpha \geq \frac{\tau_2}{\lambda} \stackrel{\text{def}}{=} \frac{1}{\lambda} \frac{V(q_1^*(\underline{\theta}), q_2^*(\underline{\theta}) \mid \underline{\theta}) - V(q_1^*(\overline{\theta}), q_2^M(\underline{\theta}) \mid \underline{\theta})}{\left(\overline{\theta} - \underline{\theta}\right)}. \tag{14}$$

See Appendix 3.2 for details. We summarise with the following proposition.

Proposition 7 *If reputational concerns are sufficiently high, 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 unobserved quality (which is lower than the unobservable quality maximising the non-reputational payoff). Formally, for* $\alpha \geq \tau_2/\lambda$, *we have* $q_1^E\left(\overline{\theta}\right) = q_1^*\left(\overline{\theta}\right)$, $q_2^E\left(\overline{\theta}\right) = q_2^*\left(\overline{\theta}\right)$ *and* $q_1^E\left(\underline{\theta}\right) = q_1^*\left(\overline{\theta}\right)$, $q_2^E\left(\underline{\theta}\right) = q_2^M\left(\underline{\theta}\right) < q_2^*\left(\underline{\theta}\right)$. *Doctors do not enjoy a reputational gain neither a loss.*

In line with the main analysis in Section 2, a *semi-separating* equilibrium arises for intermediate reputational concerns, i.e. $\tau_1 \leq \alpha \leq \tau_2/\lambda$. Under the semi-separating equilibrium the good doctor always provides high qualities $q_1^*\left(\overline{\theta}\right)$ and $q_2^*\left(\overline{\theta}\right)$. Instead, the bad doctor provides qualities $\left\{q_1^*\left(\overline{\theta}\right), q_2^M\left(\underline{\theta}\right)\right\}$ with probability $r^E = 1 - \frac{\lambda}{(1-\lambda)}\left(\frac{\alpha - \tau_2}{\tau_2}\right)$, and $\left\{q_1^*\left(\underline{\theta}\right), q_2^*\left(\underline{\theta}\right)\right\}$ with probability (r^E). Therefore, the semi-separating equilibrium smoothly connects the separating with the pooling equilibrium (see Appendix 3.3 for some additional details). Higher reputational concerns increase the probability of the bad doctor to provide the high unobservable quality and the low unobservable quality.

We conclude with two remarks which emphasise the differences between the main model and the extension with multitasking. First, in qualitative terms, the presence of multitasking increases the scope for the pooling equilibrium to arise and for policies which

higher under a pooling equilibrium since qualities are substitutes in costs and quality 1 is higher under a pooling equilibrium.

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 and saving costs on the unobservable quality. At the same time it is precisely this offsetting behaviour which makes publishing information less desirable. There may be scenarios where if patient's benefit in the unobservable quality dimension is more important than the benefit in the observable quality dimension then publishing information may be harmful to the patients. Formally, the comparison depends on $B(q_1^*(\overline{\theta}), q_2^M(\underline{\theta})) \gtrless B(q_1^*(\underline{\theta}), q_2^*(\underline{\theta}))$.

Second, the introduction of a pay-for-performance scheme (or an increase in price of an existing P4P scheme) may exacerbate the multitasking problem when reputational concerns are high. In the presence of a pooling equilibrium an increase in price will further increase the observable quality but will further decrease the unobservable quality. This also arises under a separating equilibrium. However, the 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.

5 Optimal contracting

In this section we relax the assumption that the doctor's contract is exogenous. We therefore take the analysis one step further and allow the purchaser of health services (e.g. a public or private insurer) to design the optimal contract. Formally, we endogenise 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 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 (see Rosenthal et al., 2004, and examples given in 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 possible reputational concerns over which pooling arises when evaluated at the optimal price. We proceed in a similar way for deriving the separating equilibrium.

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

The purchaser problem is

$$\underset{T,p}{Max}\; W\left(q^*\left(\overline{\theta},p\right)\right)-T-pq^*\left(\overline{\theta},p\right). \tag{15}$$

We assume that the purchaser takes into account two types of constraints: the participation constraints (PC) which ensure that each doctor is willing to provide services no matter his type, and the limited-liability (L) constraints, which ensure that each doctor does not make a negative pecuniary profit no matter his type. Formally, the PC constraints 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 L constraint is the same for both types: $T+pq^*\left(\overline{\theta},p\right)-C(q^*\left(\overline{\theta},p\right)) \geq 0$. 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. The PC constraint is always satisfied when the L constraint is since the utility is the sum of profits and the altruistic component.[8] Therefore, the L constraint is the only binding one.

Substituting for $T=C(q^*\left(\overline{\theta},p\right))-pq^*\left(\overline{\theta},p\right)$ into (15), and maximising with respect to the optimal price, we obtain:

$$\left[W_q\left(q^*\left(\overline{\theta},p\right)\right)-C_q\left(q^*\left(\overline{\theta},p\right)\right)\right]\frac{\partial q^*\left(\overline{\theta},p\right)}{\partial p}=0. \tag{16}$$

The condition simply 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, since W_q is non-increasing and C_q is decreasing, is satisfied for a unique value for $q^*\left(\overline{\theta},p\right)$, which we refer to as q^o.

The optimal contract, denoted by the pair $\left\{p^P,T^P\right\}$ is characterized by:

$$p^P = \left(1-\overline{\theta}\right)W_q\left(q^o\right), \tag{17}$$

$$T^P = C(q^o)-p^Pq^o, \tag{18}$$

and the optimal price decreases in the degree of altruism of the good doctor. In summary, under a pooling equilibrium, the purchaser can obtain allocative efficiency by designing a contract which is aimed at the good doctor. Since the bad doctor mimics the good doctor, the same contract induces allocative efficiency also for the bad doctor. This pooling equilibrium arises for $\alpha \geq \tau(p^P)/\lambda$.

Separating equilibrium. Under a separating equilibrium with low reputational concerns, the purchaser problem is:

$$\underset{T,p}{Max}\;\lambda\left[W\left(q^*\left(\overline{\theta},p\right)\right)-T-pq^*\left(\overline{\theta},p\right)\right]+(1-\lambda)\left[W\left(q^*\left(\underline{\theta},p\right)\right)-T-pq^*\left(\underline{\theta},p\right)\right] \tag{19}$$

[8] Formally, $V(q^*(\overline{\theta},p),p,T \mid \underline{\theta})=\underline{\theta}W(q^*(\overline{\theta},p))$ and $V(q^*(\overline{\theta},p),p,T \mid \overline{\theta})=\overline{\theta}W(q^*(\overline{\theta},p))$.

subject to the PC constraints: $V(q^*(\overline{\theta},p),p,T \mid \overline{\theta})+\alpha(\bar{\theta}-E(\theta))\geq 0$ 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: $T+pq^*\left(\overline{\theta},p\right)\geq C(q^*\left(\overline{\theta},p\right))$ and $T+pq^*\left(\underline{\theta},p\right)\geq C(q^*\left(\underline{\theta},p\right))$. Since profit is decreasing in altruism,[9] it is the liability constraint of the most altruistic provider which is binding, so that $T=C(q^*\left(\overline{\theta},p\right))-pq^*\left(\overline{\theta},p\right)$. The PC of the good and the bad doctor is never binding (see Appendix 4).

Substituting for T in (19) and maximising with respect to price, we obtain

$$\begin{aligned}&\lambda W_q\left(q^*\left(\overline{\theta},p\right)\right)\frac{\partial q^*(\overline{\theta})}{\partial p}+(1-\lambda)\,W_q\left(q^*\left(\underline{\theta},p\right)\right)\frac{\partial q^*\left(\underline{\theta},p\right)}{\partial p} \\ =\;& C_q(q^*\left(\overline{\theta},p\right))\frac{\partial q^*(\overline{\theta})}{\partial p}-(1-\lambda)\left[q^*\left(\overline{\theta},p\right)-q^*\left(\underline{\theta},p\right)+p\left(\frac{\partial q^*(\overline{\theta})}{\partial p}-\frac{\partial q^*\left(\underline{\theta},p\right)}{\partial p}\right)\right].\end{aligned} \tag{20}$$

The optimal price is set such that the average marginal benefit (weighted by response of quality to price) is equal to the marginal cost. The marginal cost has two components: since the L 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 and $T^S=C(q^*\left(\overline{\theta},p^S\right))-p^Sq^*\left(\overline{\theta},p^S\right)$. 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 types make 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.

We summarise in the following proposition.

Proposition 8 *If the purchaser is constrained to the use of the linear contract $\{T,p\}$ which is independent of doctor's type, higher reputational concerns make the purchaser*

[9] 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).

(weakly) better off.

This key result also holds more generally when comparing the pooling equilibrium with a semi-separating one. 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 $q^*\left(\underline{\theta}, p\right)$ 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 $T + pq^*\left(\underline{\theta}, p\right) \geq C(q^*\left(\underline{\theta}, p\right))$. Again, the PC constraints are never binding. The problem becomes

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

subject to the L constraint for the good doctor $T + pq^*\left(\overline{\theta}, p\right) \geq C(q^*\left(\overline{\theta}, p\right))$. 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. In general, the optimal price under a semi-separating equilibrium is intermediate between the price under pooling and under separating (proof omitted).

Finally, if differences in responsiveness of quality to price between the good and bad doctor is small, then it is straightforward to show that 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 prices helps to reduce the informational rent of the bad doctor.

The key policy insight is that policies aimed at enhancing reputational concerns and policies based on pay-for-performance schemes tend to be substitutes. Policymakers can induce doctors to provide higher quality either by publishing performance indicators and disseminating them widely or tie the performance indicators to specific monetary incentives. This section has shown that if the purchaser is constrained to a linear contract which is common across types, the purchaser may do better by widely disseminating information than by strengthening pay-for-performance schemes.[10]

[10] It can be shown that if the purchaser can implement more flexible non-linear contracts, then the purchaser can obtain the same welfare under the two policies. More precisely, this would be the case if the purchaser can implement a menu of contracts offering a different transfer in combination of different price level: $\{T(\theta), p(\theta)\}$ with $\theta = \underline{\theta}, \overline{\theta}$. Although these contracts are more flexible, they are rarely observed in practice (proof available from the authors).

6 Equilibrium selection and the intuitive criterion test

In this section we discuss two issues regarding the equilibria investigated throughout the paper. First, we discuss whether the beliefs that sustain the pooling equilibrium, which arises when reputational concerns are strong, pass the intuitive criterion test (ICT henceforth, Cho and Kreps, 1987; Fudenberg and Tirole, 1991). Second, we address whether other Perfect Bayesian Equilibria coexist with the pooling equilibrium.

6.1 Pooling and the intuitive criterion

Recall that the pooling equilibrium ($q^E(\theta) = q^*(\bar{\theta})$ for all θ) is sustained by assuming that observing any quality above $q^*(\bar{\theta})$ leads to the same beliefs as observing $q^*(\bar{\theta})$. Such (out-of-equilibrium) beliefs fail the ICT if there exists $\hat{q} > q^*(\bar{\theta})$ such that:

(i) The good doctor prefers quality $\hat{q}$ to $q^*(\bar{\theta})$ if $\hat{q}$ brings the best possible beliefs; and

(ii) The bad doctor prefers $q^*(\bar{\theta})$ to $\hat{q}$ even if such $\hat{q}$ leads to the best possible beliefs.

If (i) and (ii) are satisfied at $\hat{q}$ then "intuitive" beliefs should put zero probability on the doctor being bad after $\hat{q}$ has been observed. Notice that this would contradict our assumption on the beliefs for $q > q^*(\bar{\theta})$.

It turns out a quality $\hat{q}$ satisfying conditions (i) and (ii) exists for all $\alpha > \tau/\lambda$ if the single-crossing condition $V_{\theta q} \equiv W_q(q) > 0$ is satisfied for all q, i.e., if W is monotone. However, if we relax this assumption, this is not necessarily so.

We summarize our results in the next proposition (all proofs are in Appendix 5).

Proposition 9 *Suppose that benefit is quadratic,* $W(q) = v_1 q - v_2 q^2$*, and cost is linear,* $C(q) = cq$*, where* v_1, v_2, c *are positive parameters, and that the purchaser sets price equal to* $p^P = c(1-\bar{\theta})$*. Then there exists a threshold for* α*, given by*

$$\frac{\bar{\theta}}{\bar{\theta}-\underline{\theta}} \frac{c^2}{v_2} \frac{1}{1-\lambda} \stackrel{def}{=} \alpha^*, \tag{22}$$

such that the pooling equilibrium, if it arises, passes the ICT if $\alpha > \alpha^*$*.*

The proposition is illustrated in Figure 3 below (drawn, without loss of generality, for the numerical case $\underline{\theta} = \frac{1}{4}$, $\bar{\theta} = \frac{3}{4}$, $c = 1$, $v_1 = 5$ and $v_2 = \frac{1}{8}$). The horizontal line represents the threshold for α below which the separating equilibrium arises (τ). The hyperbola represents the threshold for α above which the pooling equilibrium arises (τ/λ). For α between this line and the lower horizontal line the semiseparating equilibrium arises. The increasing curve depicts α^*, given by (22).

[Figure 3 here]

Let us provide some intuition for this proposition. Notice that α^* is decreasing when $\underline{\theta}$ decreases, which makes it more likely that the pooling equilibrium is intuitive. If the bad doctor is not very altruistic, he tends to ignore the fact that patients' marginal benefits are negative when excessive care is provided. This implies that if society observes such high intensity of treatment then it becomes more plausible that it came from a bad doctor. Notice also that α^* decreases with $\bar{\theta}$. Intuitively, if the good doctor is very altruistic, he is very sensitive to a decrease in patient's benefits, which will occur if he has very high intensity of treatment. If society observes such an intense treatment then it becomes implausible that it came from a good doctor. Hence, for large enough $\overline{\theta}$ and low enough $\underline{\theta}$ there does not exist any intensity of treatment such that only good doctors would choose. Finally, α^* is small when λ is small, which also makes it more likely that the pooling equilibrium passes the ICT. If most doctors are bad, then the "average doctor" is also quite bad. Hence bad doctors find it very attractive to obtain a full reputational gain by pretending to be good doctors. This in turn implies that the good doctor should set very high, even excessive, intensity of treatment to avoid imitation. But, as already mentioned, the good doctor tends to be sensitive to this excess and will avoid it. Hence we arrive to the same conclusion: there does not exist any intensity of treatment such that only good doctors would choose.

6.2 Other equilibria and the Riley Outcome

As it is usually the case for signalling games, many PBEs coexist in large regions of parameter values. This is even more so for signalling games where the signal is continuous whereas the types are dichotomous. 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. Providing the full characterization of all the PBEs that arise in our model when one does not impose these restrictions would be a lengthy exercise (Jeitschko and Normann, 2012 provide a partial characterization.) Instead, 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 $q^*(\underline{\theta})$. By Bayes' Rule, the bad doctor suffers a full reputational loss whereas the good doctor enjoys a full reputational gain. The aforementioned quality $\tilde{q}$ is such that the bad doctor is indifferent between $\tilde{q}$ and his non-reputational optimum $q^*(\underline{\theta})$. In other words, the Riley Outcome constitutes the separating equilibrium that is least costly to the good doctor.

Formally, strategy $\tilde{q}$ is such that

$$\underbrace{V\left(q^*\left(\underline{\theta}\right)|\underline{\theta}\right) - \alpha(E\left(\theta\right) - \underline{\theta})}_{\text{Separating Equilibrium payoff}} = V\left(\tilde{q}|\underline{\theta}\right) + \underbrace{\alpha(\bar{\theta} - E\left(\theta\right))}_{\text{Reputational gain under most favorable beliefs}}. \tag{23}$$

It is also well-known that if the single-crossing condition holds (here, if $W(q)$ is monotone), then the Riley Outcome is a PBE and moreover is the only such equilibrium that passes the ICT. In contrast, we show now that under the assumptions listed in the previous proposition, the Riley Outcome ceases to be a PBE. We summarize our results in the next proposition.

Proposition 10 *Under the same assumptions listed in proposition 9 and for all α larger than*

$$\frac{\bar{\theta}}{\bar{\theta}-\underline{\theta}}\frac{c^2}{v_2}\left(\frac{1}{2}\sqrt{\frac{\bar{\theta}}{\underline{\theta}}}+\frac{\bar{\theta}+\underline{\theta}}{4\underline{\theta}}\right)\overset{def}{=}\alpha^{RO}, \tag{24}$$

*the Riley Outcome is **not** a PBE.*

The intuition is related to the one given in the previous subsection. First, notice that when reputational concerns are strong, 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 (say $\tilde{q}$) that avoids imitation is far from (and above of) the good doctor's non-reputational optimum, perhaps where marginal benefits are already negative. For sufficiently high reputational concerns (formally, for $\alpha > \alpha^{RO}$), intensity of treatment $\tilde{q}$ entails such a low benefit for the patient that the good doctor (who is very sensitive to patient's benefit) prefers his non-reputational optimum (where patient's benefit is much higher) even if this brings a reputational loss. Hence the good doctor deviates from his Riley-Outcome strategy.

We now compare and summarize the results obtained in Propositions 9 and 10.

6.3 Comparisons

The results when $W(q)$ is quadratic, $C(q)$ is linear and $p = p^P$, can be represented in a new version of Figure 3 (for the same parameter values, without loss of generality). This leads to Figure 4 below, where we have added the (upper) horizontal line representing the threshold α^{RO} (given in (24)) such that, for $\alpha > \alpha^{RO}$, the Riley Outcome is not a PBE.

[Figure 4 here]

These four lines determine several regions. We report the results in propositions 9 and 10 for each region in the same graph. The most favorable region for our equilibrium selection 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 intense enough reputational concerns and intermediate values for the prior probability that the doctor is good.

7 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 a cautious *Yes, it can.* Policies colloquially known as "name and shame", where 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". 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 to maximise patients benefits net of provider transfers, as long as the purchaser is constrained to adopt relatively simple contracts that involve a single schedule (i.e. 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. There are however some important *caveats*. First, virtuous imitation arises only when reputational concerns are sufficiently intense. Second, if doctors perform more than one task and only the performance in one dimension of quality is observable, virtuous imitation by the bad doctor will only occur in the observable quality, while he may perform extremely poorly in the other, non-observable, quality dimensions.

8 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.

Deci, E. L., 1975. Intrinsic motivation. New York: Plenum.

Delfgaauw, J., Dur, R., 2007. Signaling and screening of workers' motivation. Journal of Economic Behavior and Organization 62, 605–624.

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.

Ellingsen, T., Johannesson, M., 2008. Pride and prejudice: the human side of incentive theory. American Economic Review 98 (3), 990–1008.

Kellingsen, T., Johannesson, M., 2011. Conspicuous generosity. Journal of Public Economics 95 (9–10), 1131–1143.

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.

Fehr, E. Falk, A. 2002. Psychological foundations of incentives. European Economic Review 46 (4-5), 687-724.

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

Fluet, C., Garella, P., 2002. Advertising and prices as signals of quality in a regime of price rivalry. International Journal of Industrial Organization 20, 907–930.

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.

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.

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.

Mailath, G.J., Okuno-Fujiwara, M., and Postlewaite, A., 1993. Belief-Based refinements in Signalling Games, Journal of Economic Theory, 60, 241-276.

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.

Soetevent, A.R., 2011. Payment choice, image motivation and contributions to charity: evidence from a field experiment. American Economic Journal: Economic Policy 3 (1), 180–205.

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.

Appendix

(to be made available online only or available from authors)

9 Appendix 1. The model

9.1 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

$$V(q^*(\bar{\theta}) \mid \bar{\theta}) + \alpha(\bar{\theta} - E(\theta)) \geq \begin{cases} V(q \mid \bar{\theta}) - \alpha(E(\theta) - \underline{\theta}) \text{ for all } q < q^*(\bar{\theta}) \\ V(q \mid \bar{\theta}) + \alpha(\bar{\theta} - E(\theta)) \text{ for all } q \geq q^*(\bar{\theta}), \end{cases} \tag{25}$$

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

$$V(q^*(\underline{\theta}) \mid \underline{\theta}) - \alpha(E(\theta) - \underline{\theta}) \geq \begin{cases} V(q \mid \underline{\theta}) - \alpha(E(\theta) - \underline{\theta}) \text{ if } q < q^*(\bar{\theta}) \\ V(q \mid \underline{\theta}) + \alpha(\bar{\theta} - E(\theta)) \text{ if } q \geq q^*(\bar{\theta}). \end{cases} \tag{26}$$

Notice that the IC constraint for the high type (25) is always satisfied since the high quality $q^*(\bar{\theta})$ maximizes the non-reputational payoff $V(q \mid \bar{\theta})$, and choosing this quality instead of any other quality also maximises the reputational payoff.

As for the low type, 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 $q^*(\bar{\theta})$ (i.e. the upper expression in the right hand side of (26) is always satisfied). The provider with low altruism must also be better-off by providing the low quality $q^*(\underline{\theta})$ rather than by disguising himself by providing the higher quality $q^*(\bar{\theta})$ in the attempt of gaining the reputational payoff (lower expression in the RHS of (26)). Note that the low type has no incentive to choose a quality which is strictly above the high quality $q^*(\bar{\theta})$ since it would increase costs with no additional gains (i.e. $V(q^*(\bar{\theta}) \mid \underline{\theta}) + \alpha(\bar{\theta} - \underline{\theta})$ is maximized at $q = q^*(\bar{\theta})$ conditional on $q \geq q^*(\bar{\theta})$).

9.2 Pooling equilibrium

According to our beliefs restrictions (monotonicity and pessimism) we have that any $q < q^*(\bar{\theta})$ indicates that the doctor is bad, i.e. $\lambda^S(q) = 0$ (so $\theta^s(q) = \underline{\theta}$) for any $q < q^*(\bar{\theta})$; and that any higher quality than $q^*(\bar{\theta})$ does not provide any further information, i.e., $\lambda^S(q) = \lambda$ (so $\theta^s(q) = E(\theta)$) for $q \geq q^*(\bar{\theta})$. 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:

$$V(q^*(\bar{\theta}) \quad | \quad \bar{\theta}) \geq \begin{cases} V(q \mid \bar{\theta}) \text{ for all } q \geq q^*\left(\bar{\theta}\right) \\ V(q \mid \bar{\theta}) - \alpha(E\left(\theta\right) - \underline{\theta}) \text{ for all } q < q^*\left(\bar{\theta}\right) \end{cases} \tag{27}$$

$$V(q^*(\bar{\theta}) \quad | \quad \underline{\theta}) \geq \begin{cases} V(q \mid \underline{\theta}) \text{ for all } q \geq q^*\left(\bar{\theta}\right) \\ V(q \mid \underline{\theta}) - \alpha(E\left(\theta\right) - \underline{\theta}) \text{ for all } q < q^*\left(\bar{\theta}\right) \end{cases} \tag{28}$$

The IC constraint for the provider with high altruism 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 $\alpha\left(E\left(\theta\right) - \underline{\theta}\right)$).

The IC constraint for the provider with low altruism is satisfied only if reputational concerns are sufficiently high. To see this, consider first the upper condition in (28). 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 reduces for any q in excess of $q^*\left(\bar{\theta}\right)$. Therefore, the condition reduces to $V(q^*(\bar{\theta}) \mid \underline{\theta}) \geq V(q^*\left(\underline{\theta}\right) \mid \underline{\theta}) - \alpha(E\left(\theta\right) - \underline{\theta})$, which can be re-written as $\alpha \geq \tau/\lambda$.

9.3 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 $(1 - r)$:

$$q^E\left(\underline{\theta}\right) = \begin{cases} q^*\left(\bar{\theta}\right) \text{ with probability } 1 - r, \\ q^*\left(\underline{\theta}\right) \text{ with probability } r. \end{cases} \tag{29}$$

Then the equilibrium is characterised by:

$$r = r^E \stackrel{\text{def}}{=} 1 - \frac{\lambda}{(1-\lambda)}\left(\frac{\alpha - \tau}{\tau}\right) < 1, \tag{30}$$

$$\lambda^S\left(q^*\left(\underline{\theta}\right)\right) = 0\,,\ \lambda^S\left(q^*\left(\bar{\theta}\right)\right) = \frac{\tau}{\alpha} > \lambda, \theta^s\left(q^*\left(\underline{\theta}\right)\right) = \underline{\theta}, \tag{31}$$

$$\theta^S\left(q^*\left(\underline{\theta}\right)\right) = \underline{\theta} + \lambda^S\left(q^*\left(\bar{\theta}\right)\right)\left(\bar{\theta} - \underline{\theta}\right) = \underline{\theta} + \frac{\tau}{\alpha}\left(\bar{\theta} - \underline{\theta}\right) > \underline{\theta} + \lambda\left(\bar{\theta} - \underline{\theta}\right) = E\left(\theta\right). \tag{32}$$

Given the posited strategies, Bayes' Rule can always be applied to $q \in \{q^*\left(\bar{\theta}\right), q^*\left(\underline{\theta}\right)\}$. Posterior beliefs when either of these two qualities is observed are $\lambda^S\left(q^*\left(\underline{\theta}\right)\right) = 0$ and

$$\lambda^S\left(q^*\left(\bar{\theta}\right)\right) = \Pr\left(\theta = \bar{\theta} | q = q^*\left(\bar{\theta}\right)\right) = \tag{33}$$
$$\frac{\Pr\left(q = q^*\left(\bar{\theta}\right)|\theta = \bar{\theta}\right)\Pr\left(\theta = \bar{\theta}\right)}{\Pr\left(q = q^*\left(\bar{\theta}\right)|\theta = \bar{\theta}\right)\Pr\left(\theta = \bar{\theta}\right) + \Pr\left(q = q^*\left(\bar{\theta}\right)|\theta = \underline{\theta}\right)\Pr\left(\theta = \underline{\theta}\right)} =$$
$$\frac{\lambda}{1 - r\left(1 - \lambda\right)} > \lambda.$$

These beliefs yield the following expected types:

$$\theta^S\left(q^*\left(\underline{\theta}\right)\right) = \lambda^S\left(q^*\left(\underline{\theta}\right)\right)\bar{\theta} + \left(1 - \lambda^S(q^*\left(\underline{\theta}\right))\right)\underline{\theta} = \underline{\theta}; \tag{34}$$
$$\theta^S\left(q^*\left(\bar{\theta}\right)\right) = \lambda^S\left(q^*\left(\bar{\theta}\right)\right)\bar{\theta} + \left(1 - \lambda^S\left(q^*\left(\bar{\theta}\right)\right)\right)\underline{\theta} = \frac{\lambda\bar{\theta} + (1-r)(1-\lambda)\underline{\theta}}{1 - r(1-\lambda)}.$$

Any quality $q \notin \{q^*\left(\bar{\theta}\right), q^*\left(\underline{\theta}\right)\}$ is out of equilibrium. According to our restriction on out of equilibrium beliefs, we have that

$$\lambda^S\left(q\right) = \begin{cases} 0 \text{ for all } q < q^*\left(\bar{\theta}\right) \\ \frac{\lambda}{1-r(1-\lambda)} \text{ for all } q \geq q^*\left(\bar{\theta}\right) \end{cases} \tag{35}$$

Hence the expected type upon observation of such q is

$$\theta^S\left(q\right) = \begin{cases} \underline{\theta} \text{ for all } q < q^*\left(\bar{\theta}\right) \\ \frac{\lambda}{1-r(1-\lambda)}\bar{\theta} + \left(1 - \frac{\lambda}{1-r(1-\lambda)}\right)\underline{\theta} = \frac{\lambda}{1-r(1-\lambda)}\bar{\theta} + \frac{(1-\lambda)(1-r)}{1-r(1-\lambda)}\underline{\theta} \text{ for all } q \geq q^*\left(\bar{\theta}\right). \end{cases} \tag{36}$$

We can now determine the reputational payoff:

$$G\left(q\right) = \begin{cases} -\alpha\left(E\left(\theta\right) - \underline{\theta}\right) \text{ if } q < q^*\left(\bar{\theta}\right), \\ \alpha\left(\frac{\lambda\bar{\theta} + (1-r)(1-\lambda)\underline{\theta}}{1-r(1-\lambda)} - E\left(\theta\right)\right) \text{ if } q \geq q^*\left(\bar{\theta}\right). \end{cases} \tag{37}$$

For these strategies and beliefs to constitute an equilibrium we need three conditions. First, the low type has to be indifferent between $q = q^*\left(\bar{\theta}\right)$ and $q = q^*\left(\underline{\theta}\right)$; second, the low type has to (weakly) prefer any of the latter to setting $q \notin \{q^*\left(\bar{\theta}\right), q^*\left(\underline{\theta}\right)\}$; and third, the high type has to weakly prefer $q = q^*\left(\bar{\theta}\right)$ to $q \neq q^*\left(\bar{\theta}\right)$ despite the fact that a high output does not fully reveal his type. Using the fact that $q^*\left(\cdot\right)$ maximizes $V\left(q|\cdot\right)$, 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

$$V(q^*(\bar{\theta}) \mid \underline{\theta}) + \alpha\lambda\frac{\bar{\theta} - \underline{\theta}}{1 - r\left(1 - \lambda\right)} = V(q^*\left(\underline{\theta}\right) \mid \underline{\theta}), \tag{38}$$

$$V(q^*\left(\underline{\theta}\right) \mid \underline{\theta}) \geq \begin{cases} V(q \mid \underline{\theta}) \text{ for all } q < q^*\left(\bar{\theta}\right) \\ V(q \mid \underline{\theta}) + \alpha\left(\lambda\frac{\bar{\theta}-\underline{\theta}}{1-r(1-\lambda)}\right) \text{ for } q > q^*\left(\bar{\theta}\right), \end{cases} \tag{39}$$

$$V(q^*(\bar{\theta}) \mid \bar{\theta}) \geq \begin{cases} V(q \mid \bar{\theta}) - \alpha\lambda\frac{\bar{\theta}-\underline{\theta}}{1-r(1-\lambda)} \text{ for all } q < q^*\left(\bar{\theta}\right) \\ V(q \mid \bar{\theta}) \text{ for all } q \geq q^*\left(\bar{\theta}\right). \end{cases} \tag{40}$$

Notice that the upper expression in (39) is always satisfied because $q^*(\underline{\theta})$ maximises $V(q \mid \underline{\theta})$. The lower expression in (39) can be rewritten, using (38), as

$$V(q^*\left(\underline{\theta}\right) \mid \underline{\theta}) \geq V(q \mid \underline{\theta}) + \left\{V(q^*\left(\underline{\theta}\right) \mid \underline{\theta}) - V(q^*(\bar{\theta}) \mid \underline{\theta})\right\} \text{ for } q > q^*\left(\bar{\theta}\right), \tag{41}$$

which simplifies to

$$V(q^*(\bar{\theta}) \mid \underline{\theta}) \geq V(q \mid \underline{\theta}) \text{ for } q > q^*\left(\bar{\theta}\right). \tag{42}$$

This condition is again satisfied since $q^*(\underline{\theta})$ maximises $V(q \mid \underline{\theta})$.

Similarly, (40) 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 (38) is restrictive and equivalent to

$$\frac{\alpha\lambda}{1-r\left(1-\lambda\right)} = \frac{V(q^*\left(\underline{\theta}\right) \mid \underline{\theta}) - V(q^*(\bar{\theta}) \mid \underline{\theta})}{\bar{\theta}-\underline{\theta}} \equiv \tau \tag{43}$$

If we subtract 1 from both side we can rewrite this as

$$r = 1 - \frac{\lambda}{(1-\lambda)}\left(\frac{\alpha-\tau}{\tau}\right), \tag{44}$$

which is the equilibrium strategy provided in (8) and denoted by r^E. Substituting this expression into the expression for $\lambda^S\left(q^*\left(\bar{\theta}\right)\right)$, or (33), we obtain (after some algebra):

$$\lambda^S\left(q^*\left(\bar{\theta}\right)\right) = \frac{\lambda}{1-\left(1-\frac{\lambda}{(1-\lambda)}\left(\frac{\alpha-\tau}{\tau}\right)\right)(1-\lambda)} = \frac{\tau}{\alpha}, \tag{45}$$

as given in the main text. Then $\theta^S\left(q^*\left(\underline{\theta}\right)\right) = \underline{\theta} + \lambda^S\left(q^*\left(\bar{\theta}\right)\right)\left(\bar{\theta}-\underline{\theta}\right) = \underline{\theta} + \frac{\tau}{\alpha}\left(\bar{\theta}-\underline{\theta}\right)$.

Let us now calculate equilibrium payoffs. Recall that in order to sustain a mixed strategy with support $\{q^*\left(\bar{\theta}\right), q^*\left(\underline{\theta}\right)\}$, the low type must be indifferent between these two qualities. 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 given by $V(q^*(\bar{\theta}) \mid \bar{\theta}) + \alpha\left[\theta^S\left(q^*\left(\bar{\theta}\right)\right) - E\left(\theta\right)\right]$. Substituting $\theta^S\left(q^*\left(\underline{\theta}\right)\right) = \underline{\theta} + \frac{\tau}{\alpha}\left(\bar{\theta}-\underline{\theta}\right)$ we obtain, respectively:

$$\begin{aligned} V(q^*\left(\underline{\theta}\right) \mid \underline{\theta}) - \alpha(E\left(\theta\right) - \underline{\theta}) &= V(q^*\left(\bar{\theta}\right) \mid \underline{\theta}) + \alpha\left[\theta^S\left(q^*\left(\bar{\theta}\right)\right) - E\left(\theta\right)\right] \\ &= V(q^*\left(\bar{\theta}\right) \mid \underline{\theta}) + (\tau - \alpha\lambda)\left(\bar{\theta}-\underline{\theta}\right) \end{aligned} \tag{46}$$

$$V(q^*(\bar{\theta}) \mid \bar{\theta}) + \alpha\left[\theta^S\left(q^*\left(\bar{\theta}\right)\right) - E\left(\theta\right)\right] = V(q^*(\bar{\theta}) \mid \bar{\theta}) + (\tau - \alpha\lambda)\left(\bar{\theta}-\underline{\theta}\right). \tag{47}$$

Type $\bar{\theta}$'s payoff $V(q^*(\bar{\theta}) \mid \bar{\theta}) + (\tau-\alpha\lambda)\left(\bar{\theta}-\underline{\theta}\right)$ tends to $V(q^*(\bar{\theta}) \mid \bar{\theta}) + \tau\left(1-\lambda\right)\left(\bar{\theta}-\underline{\theta}\right)$ when α tends to τ, which is the same as the separating payoff at $\alpha = \tau$ (see (4)). Also,

$V(q^*(\bar{\theta}) \mid \bar{\theta}) + (\tau - \alpha\lambda)\left(\bar{\theta} - \underline{\theta}\right)$ 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. As with the high type payoffs, the semi-separating equilibrium tends towards the separating equilibrium when reputational concerns α tend to the lower bound τ and tends towards the pooling equilibrium where both types provide the high quality $q^*\left(\bar{\theta}\right)$ when reputational concerns α tend to the the upper bound τ/λ. This can be easily checked by inspection.

10 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 low type when the optimal quality $q^*(\underline{\theta})$ is chosen. In contrast

$$\frac{\partial V(q^*(\bar{\theta}), \underline{\theta})}{\partial p} = q^*(\bar{\theta}) + [p + \underline{\theta} W_q(q^*(\bar{\theta})) - C_q(q^*(\bar{\theta}))]\frac{\partial q^*(\bar{\theta})}{\partial p} \gtrless 0. \tag{48}$$

Higher prices increase revenues but also increase the quality of the high type, which makes it more costly for the low type to disguise as the high type. 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 high type 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), i.e. $\frac{\partial^2 q^*(\theta)}{\partial^2 \theta} < (>) 0$, 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 $\frac{\partial^2 q^*(\theta)}{\partial \theta^2}$.

11 Appendix 3. Multitasking

The problem is well behaved. The Second Order Conditions 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 $\{\underline{\theta}; \bar{\theta}\}$, $q_1^E(\theta), q_2^E(\theta)$ maximizes $\pi(q_1, q_2) + \theta W(q_1, q_2) + G(q_1)$ 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.

11.1 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 the IC constraints for both types to be satisfied. The good doctor has no incentive to mimic the bad doctor:

$$V(q_1^*(\bar{\theta}), q_2^*(\bar{\theta}) \mid \bar{\theta}) + \alpha(\bar{\theta} - E\left(\theta\right)) \geq \begin{cases} V(q_1, q_2 \mid \bar{\theta}) - \alpha(E\left(\theta\right) - \underline{\theta}) \text{ for all } q_1 < q_1^*\left(\bar{\theta}\right) \\ V(q_1, q_2 \mid \bar{\theta}) + \alpha(\bar{\theta} - E\left(\theta\right)) \text{ for all } q_1 \geq q_1^*\left(\bar{\theta}\right), \end{cases} \tag{49}$$

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

$$V(q_1^*(\underline{\theta}), q_2^*(\underline{\theta}) \mid \underline{\theta}) - \alpha(E\left(\theta\right) - \underline{\theta}) \geq \begin{cases} V(q_1, q_2 \mid \underline{\theta}) - \alpha(E\left(\theta\right) - \underline{\theta}) \text{ if } q_1 < q_1^*\left(\bar{\theta}\right) \\ V(q_1, q_2 \mid \underline{\theta}) + \alpha(\bar{\theta} - E\left(\theta\right)) \text{ if } q_1 \geq q_1^*\left(\bar{\theta}\right). \end{cases} \tag{50}$$

The latter can be re-written as: $0 \leq \alpha \leq \tau_1 \stackrel{\text{def}}{=} \frac{V(q_1^*(\underline{\theta}), q_2^*(\underline{\theta}) \mid \underline{\theta}) - V(q_1^*\left(\bar{\theta}\right), q_2^*\left(\bar{\theta}\right) \mid \underline{\theta}))}{\bar{\theta} - \underline{\theta}}$.

11.2 Pooling equilibrium

Since doctors provide the same quality $q_1^E\left(\theta\right)$, patients and society cannot distinguish between good and bad doctors. There is therefore no updating in beliefs after observing the high quality $q_1^*\left(\bar{\theta}\right)$. Hence $\lambda^S\left(q_1^*\left(\bar{\theta}\right)\right) = \lambda$ and the expected type conditional on patients observing the high quality is the average type, $\theta^s\left(q_1^*\left(\bar{\theta}\right)\right) = E\left(\theta\right)$. 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 $\lambda^S\left(q_1\right) = 0$ (so $\theta^s\left(q_1\right) = \underline{\theta}$) for any $q_1 < q_1^*\left(\bar{\theta}\right)$; and that any higher quality than $q_1^*\left(\bar{\theta}\right)$ does not provide any further information, ie $\lambda^S\left(q_1\right) = \lambda$ (so $\theta^s\left(q_1\right) = E\left(\theta\right)$) for $q_1 \geq q_1^*\left(\bar{\theta}\right)$.

For the posited strategies and beliefs to constitute a pooling equilibrium, we need again the incentive-compatibility constraints to be satisfied, so that both types of doctor have an incentive to provide the high quality:

$$V(q_1^*(\bar{\theta}), q_2^*(\bar{\theta}) \quad \mid \quad \bar{\theta}) \geq \begin{cases} V(q_1, q_2 \mid \bar{\theta}) \text{ for all } q_1 \geq q_1^*\left(\bar{\theta}\right) \\ V(q_1, q_2 \mid \bar{\theta}) - \alpha(E\left(\theta\right) - \underline{\theta}) \text{ for all } q_1 < q_1^*\left(\bar{\theta}\right) \end{cases} \tag{51}$$

$$V(q_1^*(\bar{\theta}), q_2^M\left(\underline{\theta}\right) \quad \mid \quad \underline{\theta}) \geq \begin{cases} V(q_1, q_2 \mid \underline{\theta}) \text{ for all } q_1 \geq q_1^*\left(\bar{\theta}\right) \\ V(q_1, q_2 \mid \underline{\theta}) - \alpha(E\left(\theta\right) - \underline{\theta}) \text{ for all } q_1 < q_1^*\left(\bar{\theta}\right) \end{cases} \tag{52}$$

The IC constraint for the provider with high altruism is always satisfied since (i) the non-reputational payoff $V(q_1, q_2 \mid \bar{\theta})$ is maximized at the high quality $q_1^*(\bar{\theta}), q_2^*(\bar{\theta})$ and, (ii) any other quality below $q_1^*(\bar{\theta})$ brings a reputational loss (equal to $-\alpha (E(\theta) - \underline{\theta})$).

The IC constraint for the provider with low altruism is satisfied only if reputational concerns are sufficiently high. To see this, consider first the upper condition in (28). Since the non-reputational payoff $V(q_1, q_2 \mid \underline{\theta})$ is maximized at the low quality $q_1^*(\underline{\theta}), q_2^*(\underline{\theta})$, it reduces for any q_1 in excess of $q_1^*\left(\bar{\theta}\right)$. Therefore, the condition reduces to

$$V(q_1^*(\bar{\theta}), q_2^M(\underline{\theta}) \mid \underline{\theta}) \geq V(q_1^*(\underline{\theta}), q_2^*(\underline{\theta}) \mid \underline{\theta}) - \alpha(E(\theta) - \underline{\theta}), \tag{53}$$

which can be re-written as $\alpha \geq \frac{\tau_2}{\lambda} = \frac{V(q_1^*(\underline{\theta}), q_2^*(\underline{\theta}) \mid \underline{\theta}) - V(q_1^*(\bar{\theta}), q_2^M(\underline{\theta}) \mid \underline{\theta})}{\lambda\left(\bar{\theta} - \underline{\theta}\right)}$.

11.3 Semi-separating equilibrium

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

$$r^E = 1 - \frac{\lambda}{(1-\lambda)}\left(\frac{\alpha - \tau_2}{\tau_2}\right),\ \lambda^S\left(q_1^*\left(\bar{\theta}\right)\right) = \frac{\tau_2}{\alpha}. \tag{54}$$

Then $\theta^S(q_1^*(\underline{\theta})) = \underline{\theta} + \frac{\tau_2}{\alpha}\left(\bar{\theta} - \underline{\theta}\right)$. The payoff when choosing $q^*(\underline{\theta})$, which reveals that the type is low, is given by $V(q^*(\underline{\theta}) \mid \underline{\theta}) - \alpha(E(\theta) - \underline{\theta})$. The high type's payoff is $V(q^*(\bar{\theta}) \mid \bar{\theta}) + (\tau - \alpha\lambda)\left(\bar{\theta} - \underline{\theta}\right)$.

12 Appendix 4. Optimal contracting

Separating equilibrium. 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, i.e. at $\alpha = \tau$, we obtain: $V(q^*(\underline{\theta}) \mid \underline{\theta}) - \frac{\left[V(q^*(\underline{\theta}) \mid \underline{\theta}) - V(q^*\left(\bar{\theta}\right) \mid \underline{\theta})\right)}{\bar{\theta} - \underline{\theta}}(E(\theta) - \underline{\theta}) = V(q^*(\underline{\theta}) \mid \underline{\theta})\frac{\bar{\theta} - E(\theta)}{\bar{\theta} - \underline{\theta}} + V(q^*\left(\bar{\theta}\right) \mid \underline{\theta}))\frac{E(\theta) - \underline{\theta}}{\bar{\theta} - \underline{\theta}} > 0$.

13 Appendix 5. Proofs of propositions 9 and 10

Proof of Proposition 9

Let us express the ICT formally first. Beliefs fail the ICT if there exists some $\hat{q}$ such that the following conditions simultaneously hold (recall that in the pooling equilibrium there is neither reputation gain or loss):

$$\underbrace{V\left(q^{*}\left(\bar{\theta}\right)|\bar{\theta}\right)}_{\text{eq. payoff}} < V\left(\hat{q}|\bar{\theta}\right) + \underbrace{\alpha(\bar{\theta} - E\left(\theta\right))}_{\text{Most fav. beliefs}}, \tag{55}$$

$$\underbrace{V\left(q^{*}\left(\bar{\theta}\right)|\underline{\theta}\right)}_{\text{eq. payoff}} > V\left(\hat{q}|\underline{\theta}\right) + \underbrace{\alpha(\bar{\theta} - E\left(\theta\right))}_{\text{Most fav. beliefs}}. \tag{56}$$

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

$$\tau/\lambda = \frac{1}{\lambda}\frac{1}{4}\frac{c^{2}}{v_{2}}\frac{\overline{\theta} - \underline{\theta}}{\underline{\theta}}.$$

We now express conditions (55) and (56) for the assumed functional forms.

Condition (i). At price p^{P}, type $\bar{\theta}$ obtains non-reputational payoff

$$V\left(q^{*}\left(\bar{\theta}\right)|\bar{\theta}\right) = \tag{57}$$
$$T + c\left(1 - \bar{\theta}\right)\frac{v_{1} - c}{2v_{2}} + \bar{\theta}\left(v_{1}\frac{v_{1} - c}{2v_{2}} - v_{2}\left(\frac{v_{1} - c}{2v_{2}}\right)^{2}\right) - c\frac{v_{1} - c}{2v_{2}}.$$

This type's non-reputational payoff at any other $q > q^{o}$ becomes

$$V\left(q|\bar{\theta}\right) = T + c\left(1 - \bar{\theta}\right)q + \bar{\theta}v_{1}q - \bar{\theta}v_{2}q^{2} - cq. \tag{58}$$

By inspection of the (57), the RHS of (55) decreases with q if $q > \frac{v_{1}-c}{2v_{2}} = q^{o}$, unsurprisingly. Therefore condition (i) can be expressed as $q < q^{\max}$, where $q^{\max}$ solves (55) with equality. Once expressions (57) and (58) are substituted in, the resulting equation has two solutions, but only one of them yields $q > q^{o}$, namely,

$$q^{\max} = q^{o} + \sqrt{\left(1 - \lambda\right)\alpha\frac{\overline{\theta} - \underline{\theta}}{\bar{\theta}v_{2}}}. \tag{59}$$

Condition (ii). At $p = p^{P}$, type $\underline{\theta}$ obtains the following non-reputational payoff in the pooling equilibrium:

$$V\left(q^{*}\left(\bar{\theta}\right)|\underline{\theta}\right) = V\left(q^{o}|\underline{\theta}\right) = \tag{60}$$
$$T + c\left(1 - \bar{\theta}\right)\frac{v_{1} - c}{2v_{2}} + \underline{\theta}\left(v_{1}\frac{v_{1} - c}{2v_{2}} - v_{2}\left(\frac{v_{1} - c}{2v_{2}}\right)^{2}\right) - c\frac{v_{1} - c}{2v_{2}}.$$

At any other $q > q^o$, type $\underline{\theta}$ obtains non-reputational payoff

$$V(q|\underline{\theta}) = T + c\left(1 - \bar{\theta}\right)q + \underline{\theta}v_1 q - \underline{\theta}v_2 q^2 - cq. \tag{61}$$

By inspection of the (60), the RHS of (56) 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^o. It is easy to check that $q^o > q_1$. Therefore condition (56) can be expressed as $q > q^{\min}$, where $q^{\min}$ solves (56) with equality. Once expressions (60) and (61) are substituted in, the resulting equation has again two solutions but only one of them yields $q > q^o$. Namely,

$$q^{\min} = \frac{v_1 - \frac{\bar{\theta}}{\underline{\theta}}c}{2v_2} + \sqrt{\frac{c^2}{4}\left(\frac{\bar{\theta} - \underline{\theta}}{\underline{\theta}v_2}\right)^2 + (1-\lambda)\,\alpha\frac{\bar{\theta} - \underline{\theta}}{\underline{\theta}v_2}}. \tag{62}$$

We now find a necessary and sufficient condition for conditions (i) and (ii) to be compatible.

Conditions (i) and (ii) and the threshold for pooling. Conditions (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

$$\alpha > \frac{1}{4}\frac{c^2}{v_2}\frac{1}{1-\lambda} \stackrel{def}{=} \hat{\alpha}. \tag{63}$$

(2) $q^{\max} - q^{\min} = 0$ at

$$\alpha = \frac{\bar{\theta}}{\bar{\theta} - \underline{\theta}}\frac{c^2}{v_2}\frac{1}{1-\lambda} \stackrel{def}{=} \alpha^*. \tag{64}$$

Notice also that $\alpha^* > \hat{\alpha}$ since $\frac{\bar{\theta}}{\bar{\theta} - \underline{\theta}} > 1 > \frac{1}{4}$.

Therefore, conditions (i) and (ii) are incompatible for all $\alpha > \alpha^*$. This implies that the out-of-equilibrium beliefs that support the pooling equilibrium pass the ICT test for such α. This proves the proposition. *QED*

Proof of Proposition 10

We start by computing $\tilde{q}$ as defined in (23). This condition can be rewritten as

$$V\left(q^*\left(\underline{\theta}\right)|\underline{\theta}\right) = V\left(\tilde{q}|\underline{\theta}\right) + \alpha\left(\bar{\theta} - \underline{\theta}\right). \tag{65}$$

Total differentiation of (65) with respect to α shows that $\tilde{q}$ increases with α. Indeed, $\frac{\partial\tilde{q}}{\partial\alpha} = \frac{-(\bar{\theta}-\underline{\theta})}{\underline{V}_q(\tilde{q})}$, 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 once the same assumptions listed in Proposition 9 are imposed. Solving (65) for these assumptions yields

$$\tilde{q} = \frac{1}{2\underline{\theta} v_2} \left(2\sqrt{\alpha \underline{\theta} v_2 \left(\overline{\theta} - \underline{\theta} \right)} + \underline{\theta} v_1 - c\bar{\theta} \right). \tag{66}$$

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 $q^* \left(\bar{\theta} \right)$:

$$V \left(q^* \left(\bar{\theta} \right) | \bar{\theta} \right) + \underbrace{\alpha(\underline{\theta} - E \left(\theta \right))}_{\text{Worst possible reput.payoff (negative)}} \leq V \left(\tilde{q} | \underline{\theta} \right) + \underbrace{\alpha(\bar{\theta} - E \left(\theta \right)).}_{\text{Reput.gain under the Riley Outcome}} \tag{67}$$

This can be rewritten as $V \left(q^* \left(\overline{\theta} \right) | \overline{\theta} \right) - V \left(\tilde{q} | \overline{\theta} \right) \leq \alpha \left(\overline{\theta} - \underline{\theta} \right)$. (Incidentally, both sides of last expression increase with α. Indeed, $\tilde{q}' \left(\alpha \right) > 0$ and $V \left(\cdot | \overline{\theta} \right)$ 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 $\left(\overline{\theta} - \underline{\theta} \right)$). Using the functional forms for W and C given in Proposition 9 and using $p = p^P$, condition (67) can be rewritten as $\alpha < \alpha^{RO}$ where α^{RO} is given in the proposition. *QED*

FIGURES

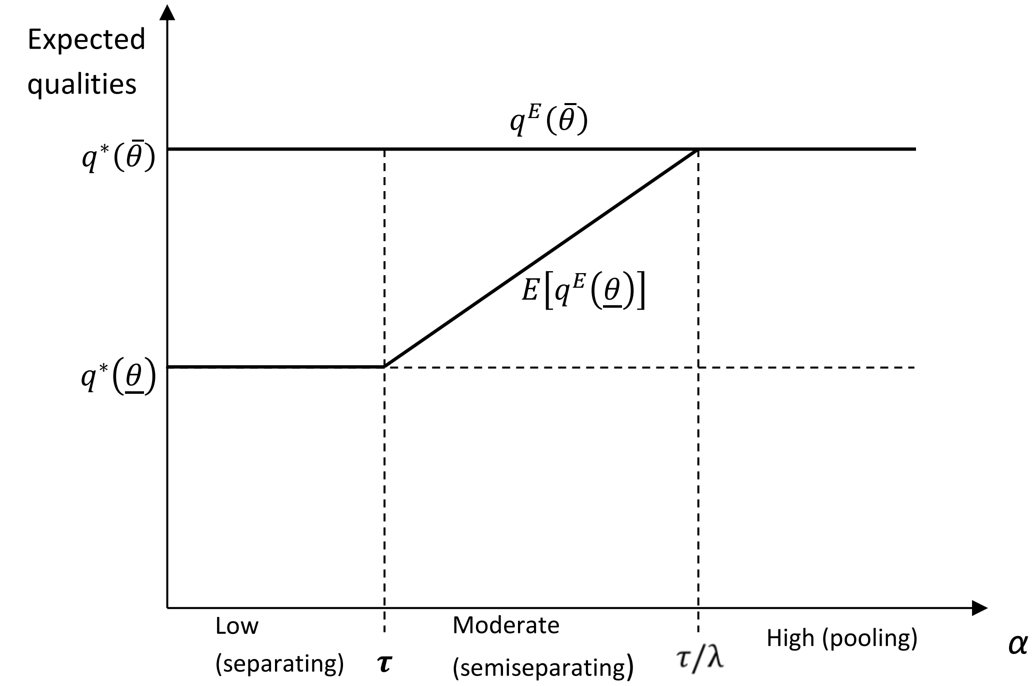

Figure 1. Expected qualities as a function of reputational concerns

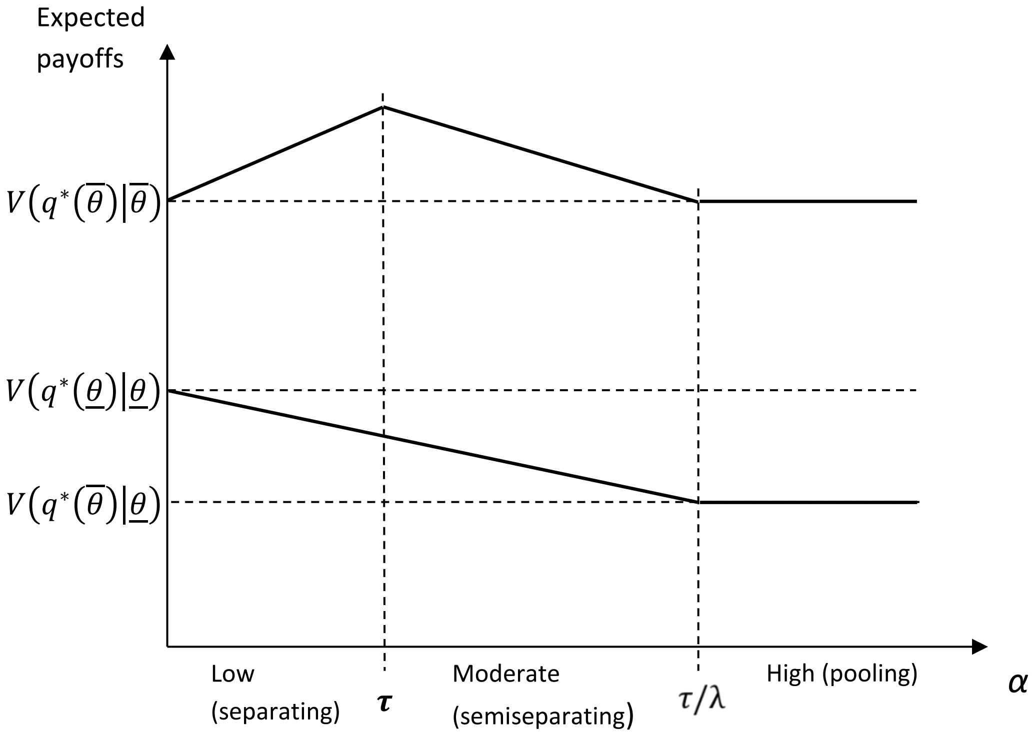

Figure 2. Expected payoffs as a function of reputational concerns

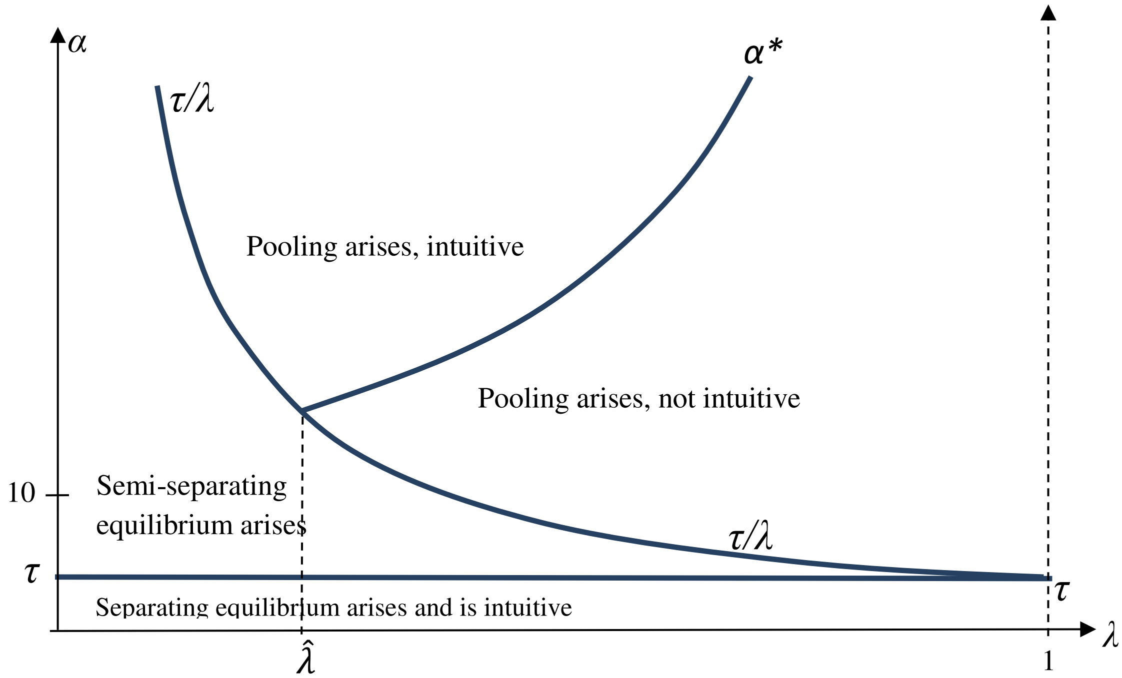

Figure 3. Equilibrium type and the Intuitive Criterion, by proportion of good doctors (λ) and Intensity of reputational concerns (*α)*

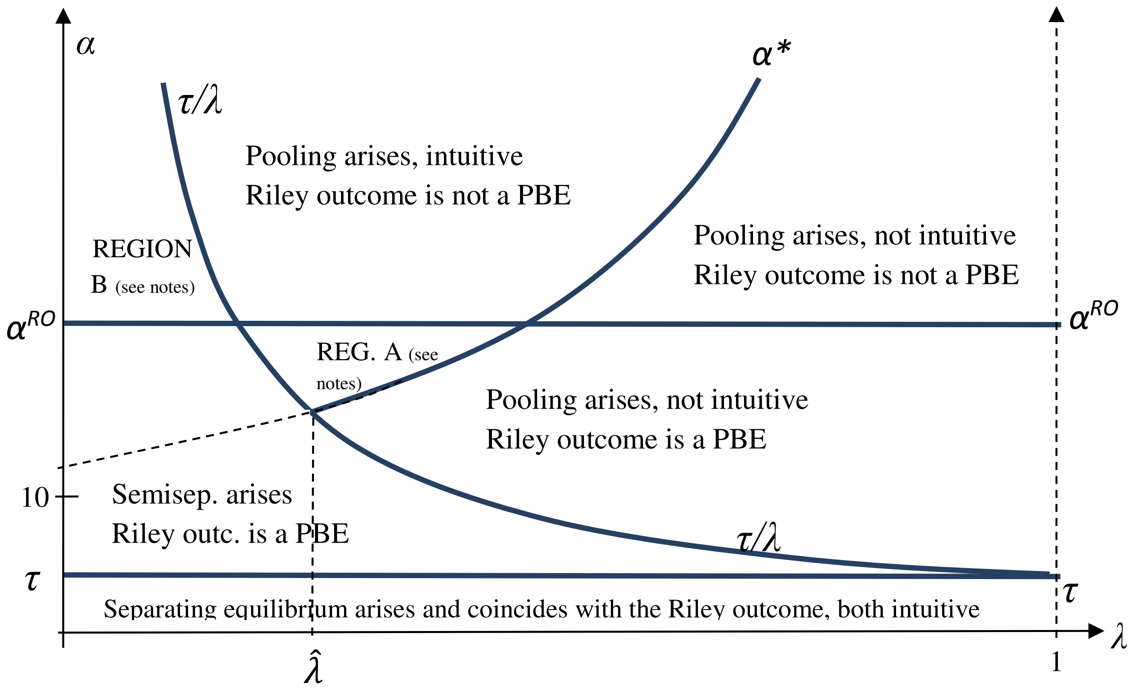

Figure 4. Equilibria and the Riley outcome, by proportion of good doctors (λ) and Intensity of reputational concerns (*α)*

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.