Truth-In-Advertising Laws and Pharmaceutical Promotion

Matthew J. Ryan and Rhema Vaithianathan

May 2008

Abstract

Prescription pharmaceuticals are an important example of a good that is purchased on the advice of a "learned intermediary" – in this case, a doctor – because of the specialised knowledge required to make an appropriate consumption decision. Direct-to-consumer advertising (DTCA) of prescription pharmaceuticals is permitted in only two OECD countries, the USA and NZ, and is highly controversial. Opponents claim that it leads to an unwarranted increase in prescribing and spiralling health care costs. Such claims often lead to calls for tighter regulation of DTCA. However, this paper suggests that tightening of regulations – more frequent auditing or increased penalties – may exacerbate rather than mitigate such problems. Because regulation enhances the credibility – and hence the profitability – of DTCA, tighter regulation can actually increase the amount of DTCA that one observes.

1 Introduction

Anti-deception laws are designed to reduce misleading advertising. In the US for example, commercial advertising is monitored by the FTC, which prosecutes firms who materially deceive consumers in a manner that is harmful. The original case for this legislation was based on an argument that it would facilitate efficient exchange by overcoming informational asymmetries (Beales, Craswell and Salop, 1981).

Some of the most stringent anti-deception laws apply to direct-to-consumer advertising (DTCA) of prescription drugs, which is monitored by the Food and Drug Administration (FDA). The FDA requires that print advertising include all important side-effects of the drug and that any television advertising include a web-site or other source of information where the side-effects can be ascertained by the patient (Vaithianathan, 2006).

Concerns have been expressed about the capability of the FDA to administer these regulations. DTCA in the US increased from $985m in 1995 to $4,237m in 2005, but the number of FDA staff responsible for monitoring DTCA has remained static. Consequently the number of violation warning letters have shown a marked decline (Donohue, Cevasco and Rosenthal, 2007).

There is surprisingly little economic analysis of truth-in-advertising regulation. In particular, we are not aware of any published paper that uses the modern theory of advertising as a quality signal to address the issue.[1] One reason for this is that in these signalling models of advertising, the content of an advertisement has no meaning – consumers respond to the fact of advertising and not its content. Indeed, Nelson (1974) argued that regulations regarding the content of advertising such as restrictions on deception are a

[1] See Pitofsky (1977), Beales, Craswell and Salop (1981) and Peltzman (1981) for older analyses from the law and economics literature. On advertising as a signal of quality, see Nelson (1974), Kihlstrom and Riordan (1984), and Milgrom and Roberts (1986).

waste of time and money. However, empirical evidence suggests otherwise. Sauer and Leffler (1990) show that following increased FTC enforcement of advertising deception, advertising intensity was noticeably affected. Surprisingly, in some sectors, improved enforcement actually *increased* advertising intensity. Sauer and Leffler conclude that stricter enforcement increased the profitability of advertising by improving its credibility.

To the extent that anti-deception laws raise the profitability of advertising, it does so for both honest and dishonest advertising. It is therefore possible that increasing the monitoring of DTCA regulation might result in more deceptive advertising. The objective of this paper is to show how such a result might occur in the context of advertising of prescription drugs.

In contrast to the earlier literature on DTCA, the present paper constructs a signalling model where the advertising need not be truthful but nonetheless reveals information in equilibrium. We analyse the effect of truth-in-advertising regulations – and their enforcement – on the prevalence of advertising and on welfare. The main conclusions are as follows: first, that improving the monitoring of regulations increases the credibility of DTCA, leading to higher levels of DTCA; and second, if consumers mistakenly believe monitoring to be more stringent than it is, then consumers might be worse off than if there were no regulation at all. The policy implications are clear: if the regulatory framework cannot be adequately monitored and prohibition of DTCA is impractical (or unconstitutional) then it might be better to completely deregulate DTCA.

We analyse a generic health system, which is meant to capture the essential elements of both a single-payer system, such as those found in New Zealand and Canada, as well as the insurance-based system of the US.[2] We

[2] New Zealand and the US are the only two OECD countries to allow DTCA of branded drugs. However, Canadian consumers are exposed to some cross-border DTCA in television broadcasts. Also, CanWest Global Communications filed a lawsuit in 2005 claiming that Canada's ban on DTCA breached the Canadian Charter of Rights and Freedoms.

assume that patients are insured and therefore do not face the full cost of prescription drugs. The physician's interests are more closely aligned to those of a Social Planner or third-party payer in the sense that they would prefer not to prescribe drugs where the cost to the third-party payer outweighs the benefits to the patient. These preferences, which are induced either through explicit supply-side cost sharing mechanisms or "softer" management approaches, create a fundamental conflict around the prescribing decision (Frank, Glazer and McGuire, 2000; Keeler, Carter and Newhouse, 1998).

Patients do not know the benefits or side effects of new drugs. Nor do they wholly trust their physicians' recommendations, suspecting them of bias against new drugs on the basis of cost. Patients might therefore apply pressure on physicians to prescribe a new drug against the physician's own recommendation. Physicians are, to some extent, persuadable through patient pressure, but not if the drug has very serious risks for the patient. In this environment, drug advertising can be a credible signal to the patient to apply pressure. Upon seeing an advertisement, the patient may infer: "the drug cannot be very bad for me, as a physician would never prescribe a harmful drug, no matter what pressure I place on her, and therefore the drug company would not advertise. The reason they advertise is to signal that if – and only if – I pressure the physician, she will prescribe the drug". In the equilibrium of our model, DTCA is used to change the doctor's prescribing behaviour in situations in which patient pressure is sufficient to induce such a change. Drugs which are clearly beneficial, and for which no pressure need be applied to elicit prescribing, will not be advertised. Nor will very harmful drugs.

The model also allows us to consider the effects of truth-in-advertising regulation. We parameterise the strictness of the regulatory environment

The case is still pending in the Ontario Supreme Court. DTCA is therefore the subject of much debate in Canada at present.

through an audit probability – which determines the probability with which misleading advertisements are detected – and a fine. We consider two alternative regulatory environments. In the first, patients are fully cognisant of the audit probability and fine. We show that increasing the audit probability or the fine makes DTCA more likely, since it increases the credibility of such advertising and is therefore welfare reducing. It may even increase the amount of *misleading* DTCA that one observes.[3]

We also consider the more plausible scenario in which patients do not know the true audit probability. We show that if consumers mistakenly believe the audit probability to be higher than it really is, then misleading DTCA may occur in equilibrium and be believed by patients. Patients may therefore end up with drugs that are potentially harmful to them. Tightening DTCA regulation – or rather, its enforcement – can have beneficial effects in such cases.

2 The consultation game

In this section, we set up a simple model of the doctor-patient consultation. We will add an advertising stage in section 4. While the model is designed for analysing the issues around DTCA, the reader will observe that its structure is quite general. It is suitable for any market in which there is a learned intermediary who advises the consumer, and whose preferences over the final decision do not perfectly match those of the consumer himself.

Consider a population of patients suffering from a particular medical condition for which two pharmaceutical therapies, α and β, are available. Drug α is a new pharmaceutical, produced under patent by firm $\mathcal{A}$. Drug β is an older, off-patent medicine, competitively supplied in generic form. Patients

[3] Patients, however, are sufficiently rational to know that the ads may be misleading, so they are not actually "mislead".

who are prescribed drug α face a larger co-payment to reflect its higher cost to a third-party payer (an insurance company or the Government).

Because patients only face a fraction of their drug costs, patients and doctors may disagree about the optimal treatment option. Doctors who are mindful of *cost-effectiveness* may lean more towards drug β than patients would like. This will be the case if the third-party payer is able to impose their preferences on the doctor through supply-side mechanisms. Because the doctor, as "learned intermediary", has superior information about the treatment options and provides advice to the patient, the patient will factor this bias into his interpretation of the advice. This creates the potential for conflict in the surgery. As we shall see, it also creates a role for DTCA to help the patient decide whether drug α really is right for them.

While our model is suitable to any situation in which the learned intermediary has a stronger inclination to recommend β than if she were a perfect agent of the patient, we shall describe an explicit scenario for concreteness. Suppose that drug β, being older, has properties that are well-known to consumers. In particular, let us assume that it has four side-effects. Drug α is newer, so its side effects are known only to doctors and the drug company through published clinical trial data. Patient uncertainty is summarised by four possible "states of the world", $\{A, a, b, B\}$. The states correspond to the number of drug α side effects revealed by the clinical trials: 0, 1, 3 and 4 side effects in states A, a, b and B respectively. Doctors and the drug company know the true state, while patients do not.

Patients are willing to bear the higher copayment for drug α provided it eliminates *at least two* side effects. In other words, a patient would prefer to be treated with α in states A and a, but prefers the old drug β in states b and B. The notation reflects the fact that drug α is much preferred in state A and slightly preferred in state a, while β is much preferable to α in state B, and slightly preferable in state b. The doctor, on the other hand, regards α as cost-effective treatment only in state A. Thus, the doctor and patient disgree about the optimal treatment in state a (and only state a). Since a Social Planner would also wish cost-effective treament decisions to be made, we shall refer to the doctor's preferred treatment regime as the "socially optimal" one.

We model the patient-doctor encounter as a *signalling game*.[4] Nature chooses the state $s \in \{A, a, b, B\}$; the doctor observes s and makes a recommendation $m \in \{“\alpha”,\ “\beta”\}$ to the patient; the patient observes the doctor's message m, but not the state s, and decides whether or not to accept the doctor's recommendation.[5] Let $\pi_s \in (0,1)$ denote the probability with which Nature chooses state s.

The doctor's recommendation is not "cheap talk" in our model. Opponents of DTCA claim that patients pressure doctors for advertised drugs and that doctors succumb to this pressure in order to avoid conflict in the surgery. It is therefore important that this potential conflict be reflected in payoffs. Let $v_s(t)$ denote the patient's utility from treatment $t \in \{\alpha, \beta\}$ in state s, and $u_s(t)$ the doctor's utility when the patient receives treatment t in state s *in accordance with the doctor's recommendation*. If the patient insists on t in state s *against the doctor's recommendation*, the doctor's payoff is $u_s(t) - c$, where $c > 0$ is a "conflict cost".

The conflict cost may be interpreted in the following way. When the doctor makes a recommendation which is immediately accepted by the patient, the length of the consultation is short. However, if the patient wants the alternative drug, the doctor has to at least make a show of sticking to his guns for fear of looking rather arbitrary in his initial recommendation. This is time consuming, and c reflects the cost of the extra time spent with the patient when the doctor's recommendation is over-turned.

[4] Calcott (1999) and De Jaegher and Jegers (2001) also use signalling models to describe the patient-doctor consultation.

[5] Since drug β is an acceptable treatment for all patients and drug α is no worse in any state, it is appropriate to accord the patient sovereignty over the final treatment choice.

The following are the basic assumptions of our model:

Assumption 1 *The utility functions satisfy:*

$$\Delta_A^v > \Delta_a^v > 0 > \Delta_b^v > \Delta_B^v$$

$$\Delta_A^u > 0 > \Delta_a^u > \Delta_b^u > \Delta_B^u$$

where

$$\Delta_s^u = u_s(\alpha) - u_s(\beta)$$

$$\Delta_s^v = v_s(\alpha) - v_s(\beta).$$

Assumption 2 *Conflict costs satisfy*

$$c < \min\{\Delta_A^u,\ |\Delta_B^u|\}.$$

Assumption 1 says that the doctor has a bias toward drug β relative to the patient's preferences. Doctor and patient rank the treatments the same way in every state except a, where the doctor prefers β and the patient α.

Assumption 2 says that the threat of conflict will be insufficient for the doctor to deviate from recommending her preferred drug in states A and B. In these states, the doctor would be happy to use "reverse psychology" if it results in her preferred drug being chosen.

3 Equilibrium in the absence of advertising

We first establish equilibrium behaviour in the absence of DTCA. Since the patient-doctor consultation is a signalling game, we employ the *perfect Bayesian Nash equilibrium (PBNE)* solution concept. We also use a standard refinement – the *intuitive criterion* of Cho and Kreps (1987) – to eliminate implausible PBNE's. Finally, from the remainder, we exclude any equilibrium that is Pareto dominated by another surviving equilibrium.[6]

The following Theorem identifies all the (surviving) equilibria of the patient-doctor signalling game.[7]

Theorem 1 *Under Assumptions 1 and 2, for any parameter values, there exists a unique Pareto undominated PBNE of the patient-doctor signalling game that satisfies the intuitive criterion. If*

$$\sum_{s\in\{a,b,B\}} \pi_s \Delta_s^v \ \leq \ 0 \tag{1}$$

it is the following: drug α is recommended in state A, drug β is recommended all other states, and the patient always follows the doctor's recommendation. If

$$\sum_{s\in\{a,b,B\}} \pi_s \Delta_s^v \ > \ 0 \tag{2}$$

it is the following:[8] *the doctor always recommends α in state A, recommends α with probability*

$$1-\left(\frac{\pi_b\left|\Delta_b^v\right|+\pi_B\left|\Delta_B^v\right|}{\pi_a\Delta_a^v}\right) \tag{3}$$

in state a, and recommends β otherwise; the patient always accepts recommendation "α", but accepts recommendation"β" with probability

$$\frac{c}{\left|\Delta_a^u\right|+c} \tag{4}$$

[6] Utility is compared from an *ex ante* perspective for the patient, and state-by-state for the doctor. (It is standard to think of each state as representing a different *'type'* of doctor.)

[7] Proofs of all Theorems may be found in the Appendix.

[8] Pitchik and Schotter (1987) study an equilibrium which is isomorphic to this one when $\pi_A = \pi_B = 0$ and $c = u_a(\alpha)$.

To understand these equilibria, suppose that the doctor were to recommend according to her preferences, suggesting "α" in state A and "β" otherwise. Under condition (1), a patient who hears the recommendation "β" and therefore learns only that $s \in \{a, b, B\}$ is happy to accept the recommendation. We therefore obtain an equilibrium in which the doctor makes her preferred recommendation in each state, and this recommendation is always accepted. In this case, equilibrium prescribing is socially optimal. If the disagreement between doctor and patient in state a arises because of *moral hazard* – α is not cost effective but the patient is insured – then this equilibrium illustrates the potential for *adverse selection* – patient ignorance of the doctor's "type" – to overcome the moral hazard problem.

However, under condition (2), the recommendation "β" will be challenged by the patient if it is known to be made in every $s \in \{a, b, B\}$. Revealing only that $s \in \{a, b, B\}$ is not enough to ensure a quiescent patient. Indeed, if this is all that the doctor reveals, then the patient will always reject recommendation "β" – not just in state a, but in states b and B also.

In equilibrium, the doctor must therefore reveal more information. This is achieved by mixing her recommendation in state a: each recommendation is made with positive probability. Effectively, the doctor gives a lukewarm recommendation of "β" in state a and a more enthusiastic recommendation in states b and B. Now, when a patient hears recommendation "β", her posterior probability on a will be

$$\frac{\theta \pi_a}{\theta \pi_a + \pi_b + \pi_B} \tag{5}$$

where

$$\theta = \frac{\pi_b \left|\Delta_b^v\right| + \pi_B \left|\Delta_B^v\right|}{\pi_a \Delta_a^v}$$

and $\theta < 1$ by condition (2). Thus,

$$\frac{\theta \pi_a}{\theta \pi_a + \pi_b + \pi_B} < \frac{\pi_a}{\pi_a + \pi_b + \pi_B}$$

and the posterior (5) is increasing in θ. In other words, as θ falls, the doctor's recommendation becomes more informative. The value of θ is low if the expected benefits of taking drug α in state a are high relative to the expected costs of taking drug α in states b or B. Thus, the lower is θ, the more inclined is the patient to reject recommendation "β" if the patient learns nothing more from such a recommendation than that the state is not A. The stonger this inclination, the more information the doctor must provide to contain the patient mutiny.

When condition (2) holds, equilibrium prescribing is not socially efficient. Neither is it patient-optimal. Drug α is prescribed with probability 1 in state A, with probability

$$1 - \left(\frac{\pi_b \left| \Delta_b^v \right| + \pi_B \left| \Delta_B^v \right|}{\pi_a \Delta_a^v} \right) \left(\frac{c}{\left| \Delta_a^u \right| + c} \right) \in (0,1) \tag{6}$$

in state a, and with probability

$$\frac{\left| \Delta_a^u \right|}{\left| \Delta_a^u \right| + c} \in (0,1)$$

in each of states b and B. For example, by choosing $\left| \Delta_a^u \right|$ sufficiently large relative to c and choosing the other parameters to make (6) sufficiently large, then we can construct an equilibrium in which drug α is sold with probability arbitrarily close to 1 in all states. More generally, when condition (2) holds, the patient is sufficiently keen on drug α that the doctor cannot use her informational advantage to prevent firm $\mathcal{A}$ from making sales outside state A.

4 DTCA

Now let's add an advertising stage to our game. To focus on the signalling role of DTCA, we shall suppose that the condition treated by drugs α and β is *chronic* – one that cannot be cured, but requires ongoing management.

Patients are therefore regularly reviewed by their doctors, removing any market-expansion incentive for DTCA.[9] We assume that all patients are initially managed with β, but at their next review the doctor must consider whether to switch them to α.

Our primary objective is to rationalise arguments by opponents of DTCA that drug advertising distorts prescribing behaviour in socially undesirable ways. Having provided such a rationalisation, we shall address the question of whether tighter regulation of DTCA alleviates or exacerbates the problem.

For this purpose, it is natural to make:

Assumption 3 *The model parameters satisfy*

$$\sum_{s\in\{a,b,B\}} \pi_s \Delta_s^v \ \leq \ 0.$$

From Theorem 1, under this assumption drug α is only prescribed in state A. Recall that this is the only state where α is cost-effective from a social welfare point of view. We wish to ascertain whether drug firm $\mathcal{A}$ can use DTCA to overcome the doctor's bias against its drug.[10] For simplicity, we shall suppose that drug β is never advertised – since it is competitively supplied, advertising expenditure could not be recovered through increased sales.

Following Nature's move, but before the doctor's, firm $\mathcal{A}$ observes the state and decides whether or not to advertise its product to consumers. Ads are observed by all. It costs \$$K$ to post an ad.

The regulation of DTCA is summarised by two parameters: the probability σ that an ad is audited by the regulator, and the fine F paid if an audited

[9] In fact, most of the drugs with significant DTCA are treatments for chronic conditions (Donohue, Cevasco and Rosenthal, 2007, p.676).

[10] As we saw in the previous section, if Assumption 3 fails it is possible that drug α has high sales in every state, so there may be little or no incentive for firm $\mathcal{A}$ to advertise anyway.

ad is found to breach the regulations. We shall suppose that ads are required to give a complete statement of the side-effects , so the content of an ad may be described by a statement of the form "s" for some $s \in \{A, a, b, B\}$. That is, the advertisment states the number of side effects that the drug imposes. An ad with content "s" posted in state s' breaches the regulations unless $s = s'$. If $s \neq s'$, then the advertising firm pays an expected fine of σF.[11]

It will be useful to define

$$\gamma \;=\; K + \sigma F.$$

The parameter γ measures the total expected cost of posting a misleading ad – it may be thought of as an index of the "credibility" of DTCA.

Finally, we define $\rho(z; q)$ to be the *change* in firm $\mathcal{A}$'s sale probability – the probability that a typical patient chooses drug α – needed to generate extra revenue of \$$z$, given current sales probability q. We shall impose the following:

Assumption 4 *For any $q \in [0, 1]$ there exists $z(q) \geq 0$ such that $\rho(z; q)$ is well defined iff $0 \leq z \leq z(q)$. Moreover, $\rho(z; q) \in [0, 1-q]$ when $0 \leq z \leq z(q)$, $\rho(0; q) = 0$ and $\rho(z; q)$ is strictly increasing in z when $0 \leq z < z(q)$.*

This simply says that higher sales probability translates into higher revenue.

A simple characterisation of $\rho(z; q)$ may be given when there exists a fixed number N of patients, each of whom generates revenue of \$$r$ as a customer of firm $\mathcal{A}$. In this case,

$$z(q) \;=\; (1-q)\, rN$$

and

$$\rho(z; q) \;=\; \frac{z}{rN}$$

[11] For simplicity, we assume that any falsehood attracts a sanction, even if firm $\mathcal{A}$ claims *more* side-effects than drug α actually possesses. This is innocuous – in equilibrium firm $\mathcal{A}$ will never have any incentive to make such a claim, even if it did not attract a fine.

whenever $z \leq z(q)$.

Finally, we shall make:

Assumption 5 *The cost K satisfies $K < z(0)$.*

This says that if firm $\mathcal{A}$ is currently making no sales, then an advertising campaign that captures the entire market and costs $\$K$ will strictly increase profit. Clearly, without such an assumption we would never observe advertising in equilibrium.

Since the game with advertising is no longer of the signalling variety, we apply the notion of *sequential equilibrium (SE).* We continue to select Pareto-undominated equilibria.[12]

Theorem 2 *Consider the following strategy profile:*

Firm $\mathcal{A}$ *advertises truthfully in state a and not at all in any other state;*

The doctor *always recommends drug α in state A, always recommends β in states b and B, recommends drug α in state a if firm $\mathcal{A}$ posts an ad with content "a", and recommends β in state a otherwise;*

The patient *always accepts recommendation "α", accepts recommendation "β" if he observes no ad or an ad with content "s"$\neq$"a", and accepts recommendation "β" with probability θ otherwise.*

Under Assumptions 1, 2, 4 and 5, there exists a Pareto undominated SE with these strategies for some $\theta \in [0,1]$ iff

$$\frac{c}{|\Delta_a^u|} \geq \frac{1 - \rho(\min\{\gamma, z(0)\}; 0)}{\rho(\min\{\gamma, z(0)\}; 0)} \tag{7}$$

[12] The drug firm's welfare is evaluated state-by-state.

In this equilibrium, firm $\mathcal{A}$ advertises in state a only. The doctor makes a patient-optimal treatment recommendation in each state, which the patient always accepts. DTCA allows the patient to overcome the asymmetric information problem.

The equilibrium possesses a straightforward logic. If there were no DTCA in state a, the doctor would switch her drug recommendation to "β", which the patient would accept, mistakenly believing the state to be b or B. This justifies firm $\mathcal{A}$'s decision to advertise in state a – recall that truthful advertising which captures the whole market is profitable (Assumption 5). But if firm $\mathcal{A}$ tried to post ad "a" in state b or B, the doctor would challenge the advertised claim, putting doubt in the patient's mind. As a result, the patient mixes, choosing β with probability θ, which must be sufficiently high to discourage misleading DTCA.[13]

The rationale for condition (7) is also clear. If

$$\frac{c}{|\Delta_a^u|}$$

were too low, the doctor would not be prepared to recommend α in state a, preferring to challenge firm $\mathcal{A}$'s DTCA and inducing the patient to randomise. The cost of doing so is the conflict endured when the patient chooses α, while a benefit accrues when the patient chooses β. The cost therefore depends positively on c, while the benefit depends positively on $|\Delta_a^u|$.

In the context of Assumption 3, Theorems 1 and 2 provide a logically coherent basis for claims that DTCA pressures doctors into over-prescribing. In the absence of DTCA, prescribing is socially optimal (Theorem 1 under Assumption 3); but when DTCA is allowed, the prescribing of the more expensive drug α might increase (Theorem 2):

Corollary 1 *Under Assumptions 1- 5, firm $\mathcal{A}$ sells drug α only in state A*

[13] Condition (7) ensures that a suitably high θ value can be found.

when DTCA is banned, but if DTCA is permitted, there is an equilibrium in which it sells drug α in state a as well, provided condition (7) holds.

5 Monitoring of DTCA

Interestingly, the social costs of DTCA need not be ameliorated by more stringent auditing of the regulations or increased fines for breaches of them. In fact, the reverse may be true. Under Assumption 4, increasing γ will (weakly) relax condition (7) and therefore expand the range of other parameter values for which the equilibrium in Theorem 2 exists. We therefore obtain:

Corollary 2 *Under Assumptions 1, 2, 4 and 5, increasing the "credibility" of DTCA – that is, increasing the value of γ – will increase the amount of DTCA that one observes. In particular, increasing σ or F will encourage more DTCA. (Note that these parameters can be increased without fear of jeopardising Assumption 5, unlike increases in K.)*

6 Misleading DTCA

The previous results show how tightening the regulation of DTCA may increase the amount of DTCA observed in state a – in other words, *truthful* DTCA. Such DTCA may nevertheless be deemed undesirable if drug α is not cost-effective in state a.

What about the potential for DTCA to *mislead* patients – a more worrisome phenomenon? Surely this will be curtailed by tighter regulation?

Not necessarily. Certainly, the tighter the regulation, the lower the incentive to make false claims. However, doctors are more likely to challenge

DTCA the more misleading it is. If patients are more receptive to unchallenged DTCA, this creates the potential for an equilibrium in which firm $\mathcal{A}$ uses misleading DTCA, but not too misleading. The patient knows that the equilibrium behaviour of the doctor will prevent him from making too grievous an error, so DTCA may be a reliable signal that it is in his interest to request drug α. Tighter regulation undermines incentives for misleading DTCA in state B – where it is most likely to be challenged – before it does so in state b, so tighter regulation may help support an equilibrium in which DTCA occurs (truthfully) in state a and (misleadingly) in state b. The following Theorem illustrates this possibility.

Theorem 3 *Consider the following strategy profile:*

Firm $\mathcal{A}$ *posts ad "a" in states a and b, but does not advertise in any other state;*

The doctor *always recommends drug α in state A, always recommends β in state B, recommends drug α in state $s \in \{a, b\}$ if firm $\mathcal{A}$ posts an ad with content "a" and recommends β in state $s \in \{a, b\}$ otherwise;*

The patient *always accepts recommendation "α", accepts recommendation "β" if he observes no ad or an ad with content "s"$\neq$"a", and accepts recommendation "β" with probability θ otherwise.*

Under Assumptions 1, 2, 4 and 5, there exists an SE with these strategies for some $\theta \in [0, 1]$ iff

$$\gamma \leq z(0) \tag{8}$$

$$\frac{c}{|\Delta_b^u|} \geq \frac{1 - \rho(\gamma; 0)}{\rho(\gamma; 0)} \tag{9}$$

and

$$\sum_{s \in \{a,b\}} \pi_s \Delta_s^v \geq 0. \tag{10}$$

In this equilibrium, drug α is prescribed in all states except B. Once again, increases in σ or F relax condition (9) and in this sense make it easier to sustain the equilibrium. Of course, such changes tighten condition (8), so the result is not quite unambiguous.[14] But for fixed K, and assuming $\gamma < z(0)$ initially, tightening of regulation will expand the range of conditions under which this equilibrium can be supported.

It is also worthy of note that regulations which punish according to the degree of inaccuracy may be particularly counter-productive in this respect. The reason the drug firm can advertise to consumers profitably in state b is that the punishment risk, combined with the increased likelihood of doctor resistance, discourages advertising in state B. Consumers are willing to accept treatment α in the presence of DTCA precisely because they know that drug α cannot be *too* bad for them. In equilibrium, the patient knows full well that DTCA may be misleading – they may be in state b – so it is only their certainty of *not* being in state B that leads them to choose the new drug. Penalties which are tailored to the degree of falsehood help support such beliefs. Firm $\mathcal{A}$ might be quite happy to accept draconian penalties for advertising in state B in order to support sales in state b which it otherwise could not make.

For later reference, we also note the following "converse" to Theorem 3:

Theorem 4 *Under Assumptions 1 and 4, if*

$$\sum_{s \in \{a,b\}} \pi_s \Delta_s^v < 0 \tag{11}$$

there is no SE in which Firm $\mathcal{A}$ posts an ad with content "a" in states a and b, and does no advertising in the other two states.

[14] We have also not established that the equilibrium is Pareto undominated. However, the main point is clear – the same sort of credibility effects of regulation that help support truthful DTCA play an analogous role in supporting misleading DTCA as well.

In all equilibria, including that of Theorem 3, patients are cynical rationalists,[15] and are never actually "mislead" – they realise that an ad from firm $\mathcal{A}$ may be received in state b. However, there is some evidence that consumers over-estimate the rigour of DTCA regulation and may therefore place more credence in advertisements than is warranted.[16]

To capture this possibility within our model, let us suppose that, in addition to choosing the state, Nature also chooses whether or not DTCA is regulated. By "regulated" we shall now mean that *all* DTCA is audited *with probability 1.* "Unregulated" DTCA is *never* audited. Once again, drug firms and doctors observe Nature's move while patients do not.

As before, Nature chooses state s with probability π_s. For any given state, the probability that Nature chooses "regulated" DTCA is σ. Thus, σ still represents the probability of an audit, but it is now a subjective assessment by patients – the actual audit probability is either zero ("unregulated") or one ("regulated"). In the extreme cases – $\sigma = 0$ or $\sigma = 1$ – the present model coincides with the previous one. For this modified game, we have the following result:

Theorem 5 *Suppose that condition (11) obtains. Consider the following strategy profile for the modified game:*

Firm $\mathcal{A}$ *posts ad "a" in state a if DTCA is regulated and posts ad "a" in states a and b if DTCA is unregulated. It does not advertise in any other contingency;*

The doctor *always recommends drug α in state A, always recommends β in state B, recommends drug α in state a or b if firm $\mathcal{A}$ posts an ad with content "a" and recommends β in state a or b otherwise;*

[15] Vaithianathan (2006, p.236) summarises survey evidence that consumers are indeed quite cynical about the motives behind DTCA.

[16] Wilkes, Bell and Kravitz (2000, p.118).

The patient *always accepts recommendation "α", accepts recommendation "β" if he observes no ad or an ad with content "s"≠"a", and accepts recommendation "β" with probability*

$$\theta = \frac{c}{|\Delta_b^u| + c}$$

otherwise.

Under Assumptions 1, 2, 4 and 5, there exists an SE with these strategies iff

$$\frac{c}{|\Delta_b^u|} \geq \frac{1-\rho(K;0)}{\rho(K;0)}$$

and σ *satisfies*

$$\sigma\Delta_a^v + (1-\sigma)\sum_{s\in\{a,b\}} \pi_s \Delta_s^v \geq 0 \tag{12}$$

Given (11), condition (12) requires that σ is sufficiently high – it is clearly satisfied when $\sigma = 1$, for example. In other words, the consumer must be sufficiently optimistic about the rigour of regulation. If DTCA is actually "unregulated", then such beliefs reflect misplaced confidence, allowing firm $\mathcal{A}$ to mislead the consumer in state b. Comparing Theorems 4 and 5, we see that if consumers over-estimate the audit probability, then increasing the frequency of audits to bring it in line with consumer expectations may eliminate misleading DTCA.

7 Concluding comments

Empirical studies have found two distinct effects of DTCA. The first effect is to change the physician's prescribing decision. This is the effect that we are concerned with in our paper. Wosinska (2002) studies panel data on cholesterol drugs from a private insurer and finds that DTCA increases demand for the advertised brand's drug – at least in the case where the drug is on the formularly. Narayanan, Desiraju and Chintagunta (2004) also find a positive effect of DTCA on own-brand prescribing.

The second effect is a *market expanding effect*, which increases demand for all drugs in the advertised class rather than just the named drug. Evidence of this market expanding effect is found by Rosenthal *et al.* (2003) and Iizuka and Jin (2005).

While our paper focusses on the prescribing-distortion effect of DTCA, we recognise that its relative significance is an empirical question that is yet to be definitively answered.[17] However, if DTCA simply increases physician visits, there would seem to be no compelling reason to regulate it. Regardless of the empirical evidence, the policy implications of this paper are useful. First of all, for better or worse, DTCA is regulated and recent calls to both tighten the rules and increase the FDA capacity to monitor these rules need to be analysed. Moreover, the conclusion reached in this paper is that even if DTCA is socially harmful, it might be better to completely deregulate it, in which case the optimal regulatory framework is independent of whether DTCA distorts prescribing or not.

Appendix

A Proof of Theorem 1

It is straightforward to verify that the strategy profiles described in Theorem 1 are PBNE's and that each is robust to the intuitive criterion. The work to be done is to exclude the existence of any other (acceptable) equilibria.

[17] In fact, our results can also be interpreted in a market-expanding context. If β were a non-pharmaceutical treatment – such as exercise or dietary changes – then the promotion of drug α would have a market expanding effect.

We may classify PBNE's according to whether the doctor's strategy calls for the *same* (possibly mixed) message to be sent in *every* state (a *pooling* equilibrium) or not (a *discriminating* equilibrium).[18] In the latter case, each message $m \in \{$"α", "β"$\}$ is sent with positive probability in equilibrium, and each message induces a *different* posterior belief in the patient.[19]

We first show that there is no Pareto-undominated pooling equilibrium that satisfies the intuitive criterion.

In any pooling equilibrium, the patient necessarily retains his prior beliefs after hearing the doctor's recommendation. We therefore distinguish three cases, depending on which treatment the patient prefers given his prior beliefs.

Case I. If

$$\sum_{s\in\{A,a,b,B\}} \pi_s \Delta_s^v \; < \; 0$$

then the patient always chooses β in response to any message sent with positive probability in a pooling equilibrium. It is easy to see that we cannot have an equilibrium in which the doctor sends message "α" in every state – she will certainly want to deviate in state A, no matter how the patient responds. Similarly, we can rule out equilibria in which the doctor strictly mixes in every state. She would only be indifferent in state s if

$$\Delta_s^u \; - \; c \; = \; 0,$$

which cannot hold for all s.

[18]With only two possible messages it is not possible to achieve full separation of all states.

[19]If the posteriors are the same, then they must coincide with the prior beliefs. Such cases are covered by our analysis of pooling equilibria. In particular, for such a scenario to arise the doctor must strictly mix in every state – recall that $\pi_s > 0$ for all s.

Suppose, then, that the doctor's strategy is to send message "β" in every state. Then we obtain a PBNE provided the patient responds to the off-equilibrium message "α" by selecting β with sufficiently high probability (eg, one) – otherwise the doctor has an incentive to change her recommendation in state A. However, this PBNE fails the intuitive criterion. In any state but A, if the doctor switches to the other message her payoff *must fall* relative to her equilibrium payoff: either the patient ends up choosing α or there is conflict. Deviation is therefore equilibrium-dominated in all states but A. Hence, in state A the doctor will deviate to "α" as the patient will accept this recommendation.

Case II. Next, consider a pooling equilibrium with

$$\sum_{s\in\{A,a,b,B\}} \pi_s \Delta_s^v \; > \; 0.$$

In this case, the patient will respond to any on-equilibrium message by choosing α. Then we cannot have an equilibrium in which the doctor sends message "β" in every state – she will certainly want to deviate in states a, b and B. As for Case I, we can rule out equilibria in which the doctor strictly mixes in every state. She would only be indifferent in state s if

$$\Delta_s^u \; + \; c \; = \; 0,$$

which cannot hold for all s.

Suppose, then, that the doctor sends message "α" in every state. We obtain an equilibrium provided the patient responds to any off-equilibrium message by choosing α with sufficiently high probability (eg, one). However, if

$$\sum_{s\in\{a,b,B\}} \pi_s \Delta_s^v \; \leq \; 0,$$

then any such pooling equilibrium is Pareto dominated by the first PBNE in Theorem 1: the doctor is better off in every state except

A, where her payoff is the same, while the patient's *ex ante* expected utility is (weakly) improved. Conversely, if

$$\sum_{s \in \{a,b,B\}} \pi_s \Delta_s^v \; > \; 0,$$

then any such pooling equilibrium is Pareto dominated by the second PBNE in Theorem 1. To see why, note that the patient chooses α with positive probability in response to either message in the Theorem 1 PBNE. It follows that α is an optimal choice given either posterior. Therefore, choosing α with probability 1 in response to either message must yield the same *ex ante* payoff. For the doctor, she is strictly better off in states b and B, while both equilibria deliver the same utility in each of states A and a.

Case III. Finally, consider a pooling equilibrium with

$$\sum_{s \in \{A,a,b,B\}} \pi_s \Delta_s^v \; = \; 0 \tag{13}$$

If any such equilibrium exists, it is certainly Pareto dominated by the first equilibrium in Theorem 1. It is obvious that the doctor cannot be worse off in any state in the Theorem 1 PBNE. To see that the patient is strictly better off, note that the patient's *ex ante* utility from a pooling equilibrium under condition (13) is the same as his *ex ante* utility from getting treatment α in all states. Since (13) implies

$$\sum_{s \in \{a,b,B\}} \pi_s \Delta_s^v \; < \; 0,$$

the patient's *ex ante* utility must be higher in the Theorem 1 PBNE: replacing treatment α with treatment with β in event $\{a, b, B\}$ strictly increases *ex ante* expected utility.

Having eliminated all the pooling equilibria, we now turn to the discriminating PBNE's. In any such equilibrium, each message is sent with positive probability in equilibrium – there are no off-equilibrium messages. Moreover, each message induces a different posterior in the patient. Without loss of generality, let us suppose that one message elicits response α with probability q and the other elicits α with probability $q' \geq q$. We shall call the former the "q message" and the latter the "q' message".

We first rule out $q' = q$.

If $q' = q = k$ and $k \neq \frac{1}{2}$, the doctor would have a *strict* preference for whichever message minimised conflict, and his preference would be the same in *every* state. This contradicts the assumption that two messages are sent in equilibrium. Suppose, then, that $k = \frac{1}{2}$. In this case, the patient is happy to choose either treatment conditional on either message. It follows that he is just as happy to choose α in response to either message, or to choose β in response to either message. This in turn implies that he is indifferent between α and β given his prior beliefs:

$$\sum_{s \in \{A,a,b,B\}} \pi_s \Delta_s^v \ = \ 0.$$

From this equation we deduce that

$$\sum_{s \in \{a,b,B\}} \pi_s \Delta_s^v \ < \ 0.$$

We now see that any such equilibrium must be Pareto dominated by the first PBNE in Theorem 1: the doctor is better off in every state and the patient is better off *ex ante*.[20]

Therefore, we conclude that $q' > q$. In particular, $q' > 0$ and $q < 1$.

It is clear from Assumption 2 that the doctor strictly prefers to send the q' message in state A and strictly prefers to send the q message in state B. It is also obvious that if the q message is sent with positive probability in state

[20]Recall that the *ex ante* utility of the current equilibrium is the same as if the patient received α in every state.

a, then it is sent with probability 1 in state b. In fact, the q message *must* be sent with positive probability in state a. If this were *not* the case, then $q = 0$ would be necessary for patient-optimal behaviour, since states A and a are ruled out by the q message. But then the doctor will definitely prefer to send the q message in state a – a contradiction.

Thus:

- *the q' message is sent with probability 1 in state A*
- *the q message is sent with probability 1 in states b and B*
- *the q message is sent with positive probability in state a*, and
- $q' = 1$ (since the q' message rules out states b and B).

This leaves two possible scenarios to consider:

Scenario I: q message sent with probability 1 in state a. A Scenario I equilibrium can only exist if

$$\sum_{s \in \{a,b,B\}} \pi_s \Delta_s^v \ \leq \ 0 \tag{14}$$

If this condition did not hold, the patient's optimal response to the q message is to choose α with probability 1 – but we know that $q < 1$. Moreover, if condition (14) holds, then the Theorem 1 equilibrium Pareto dominates any other Scenario I PBNE: if $q > 0$ the patient at least weakly prefers the Theorem 1 equilibrium while the doctor is weakly better off in all states and strictly so in (at least) states a, b and B; if $q = 0$ but the q message is "α", the patient is indifferent between equilibria but the doctor strictly prefers the Theorem 1 equilibrium in all states.

Scenario II: q **message sent with probability** $p \in (0,1)$ **in state** a**.** For the doctor to be happy to randomise in state a, we require

$$q\left[u_a(\alpha) - c\right] + (1-q)\, u_a(\beta) \;=\; u_a(\alpha) \quad \Leftrightarrow \quad q \;=\; \frac{c}{|\Delta_a^u| + c} \tag{15}$$

if the q message is "β". If, on the other hand, the q message is "α", the doctor always strictly prefers to send the q message in state a:

$$\begin{aligned} qu_a(\alpha) + (1-q)\left[u_a(\beta) - c\right] \;&>\; u_a(\alpha) - c \\ \Leftrightarrow \quad u_a(\beta) - c \;&>\; u_a(\alpha) - \frac{c}{(1-q)} \end{aligned}$$

which holds for any $q \in [0,1)$. We may therefore assume that the q message is "β" and that q satisfies (15). The latter implies $q \in (0,1)$ and hence the patient must be indifferent between treatment options after hearing "β":

$$\begin{aligned} p\pi_a\left[v_a(\alpha) - v_a(\beta)\right] - \pi_b\left[v_b(\beta) - v_b(\alpha)\right] - \pi_B\left[v_B(\beta) - v_B(\alpha)\right] \;&=\; 0 \\ \Leftrightarrow \quad p \;&=\; \frac{\pi_b\,|\Delta_b^v| + \pi_B\,|\Delta_B^v|}{\pi_a \Delta_a^v} \end{aligned} \tag{16}$$

Provided the p defined by (16) is strictly between 0 and 1, we can construct an equilibrium of this sort. Note that $p > 0$, but $p < 1$ iff

$$\sum_{s \in \{a,b,B\}} \pi_s \Delta_s^v \;>\; 0$$

So provided this condition is met – and only then – we obtain a unique PBNE in which the doctor mixes in state a. It is precisely the PBNE described in Theorem 1.

B Proof of Theorem 2

The proof proceeds as follows. First, taking $\theta \in [0,1]$ as given, we determine conditions on θ for the indicated strategy profile to constitute (part of) a sequential equilibrium. Next, we show that there exist θ values in $[0,1]$ satisfying these conditions iff (7) holds. Finally, we show that any such SE is not Pareto dominated.

Therefore, let us start by taking $\theta \in [0,1]$ as given and establishing conditions under which the stated strategies, together with some consistent beliefs for the patient, form a SE.

Drug firm. It is clearly rational for firm $\mathcal{A}$ to advertise truthfully in state a. If it withdrew its ad altogether, the doctor would switch her recommendation to "β" and the patient would accept this recommendation. By Assumption 5, this would lower firm $\mathcal{A}$'s profit. If firm $\mathcal{A}$ posted a misleading ad in state a, it would not only lose all its sales, but its advertising costs would also rise (at least weakly) from K to γ. By Assumptions 4 and 5, this would result in a loss of profit.

It is also rational for firm $\mathcal{A}$ *not* to advertise in any other state. An ad in state A, misleading or otherwise, would change neither the doctor's recommendation nor the patient's response, so is clearly a waste of money. Irrespective of its content, an ad posted in state b or B would not change the doctor's recommendation, "β". If the ad content is not "a", the patient will choose β, so advertising is again a waste of money. If the ad content is "a", the patient will choose α with probability $1-\theta$. The firm's profit will increase by such a deviation only if $1-\theta > \rho\left(\min\left\{\gamma, z\left(0\right)\right\}; 0\right)$.

Therefore, firm $\mathcal{A}$'s strategy is sequentially rational iff

$$\theta \;\geq\; 1-\rho\left(\min\left\{\gamma, z\left(0\right)\right\}; 0\right) \tag{17}$$

Doctor. It is easy to see that the doctor's strategy is sequentially rational in state A, even if firm $\mathcal{A}$ deviates from its equilibrium behavior – the patient is always happy to accept recommendation "α", no matter what firm $\mathcal{A}$ does.

Consider state $s \in \{b, B\}$. It is clearly optimal for the doctor to recommend β in all scenarios other than the one in which firm $\mathcal{A}$ posts an ad with content "a", since the patient will be happy to accept the recommendation. If firm $\mathcal{A}$ *does* post an ad with content "a", recommending β is optimal iff

$$u_s(\alpha) \leq \theta u_s(\beta) + (1-\theta)[u_s(\alpha) - c] \quad \Leftrightarrow \quad \frac{c}{|\Delta_s^u|} \leq \frac{\theta}{1-\theta} \tag{18}$$

Consider state a. Once again, it is optimal to recommend β if firm $\mathcal{A}$ does anything other than post an ad with content "a", since the patient will happily accept the recommendation. If firm $\mathcal{A}$ posts an ad with content "a", it is optimal to recommend α iff

$$u_a(\alpha) \geq \theta u_a(\beta) + (1-\theta)[u_a(\alpha) - c] \quad \Leftrightarrow \quad \frac{c}{|\Delta_a^u|} \geq \frac{\theta}{1-\theta} \tag{19}$$

Therefore, the doctor's strategy is sequentially rational iff condition (19) holds and condition (18) holds for $s \in \{b, B\}$. Since $|\Delta_b^u| < |\Delta_B^u|$, these three conditions will be satisfied iff

$$\frac{c}{|\Delta_b^u|} \leq \frac{\theta}{1-\theta} \leq \frac{c}{|\Delta_a^u|} \tag{20}$$

Patient. The patient's behaviour is clearly optimal in the absence of a detectable deviation from equilibrium by the other players. The detectable off-equilibrium scenarios may be usefully divided into three cases: (a) recommendation "β" in conjunction with an ad with content "a" by firm $\mathcal{A}$; or (b) recommendation "β" in conjunction with an ad by firm $\mathcal{A}$ with content "s"$\neq$"a"; or (c) recommendation "α" in conjunction with an ad by firm $\mathcal{A}$ with content "s"$\neq$"a". Each scenario can be rationalised by supposing that only one player has deviated. In the case of scenarios (b) and (c), only one such rationalisation is possible; while for (a) there are two. Let us therefore suppose the patient interprets (c) to imply that only firm $\mathcal{A}$ has deviated and hence that the state is A; (b) to imply that only firm $\mathcal{A}$ has deviated and hence that the state is a, b or B; and (a) to imply that either the doctor deviated and we are in state a, or that the drug firm deviated and we are in state b or B. In the latter case, let us further suppose that the patient's updated beliefs make him indifferent between treatments. With such beliefs, the patient's strategy is sequentially rational.

It remains to verify that such beliefs are *consistent*, in the sense of Kreps and Wilson (1984). The informal rationalisation of the beliefs given above strongly suggests that they are. The reader who is happy with the informal analysis may skip the following computations without loss of continuity or comprehension. For the rest, here are the formalities. To construct consistent beliefs, we must first construct a sequence of strictly mixed stategies for each player that converge to the equilibrium profile. Let us index this sequence of profiles by $k \in \{1, 2, ...\}$.[21] Let $\{\varepsilon_k\}_{k=1}^{\infty}$ be a sequence of numbers in $(0,1)$ that converges to 0 as $k \to \infty$, and satisfying

$$\left(\frac{\pi_b \left|\Delta_b^v\right| + \pi_B \left|\Delta_B^v\right|}{\pi_a \Delta_a^v}\right) \varepsilon_k \ < \ 1 \qquad (21)$$

for each k. Suppose that, in every state, the drug firm plays each of its four off-equilibrium actions with probability

$$\frac{\varepsilon_k}{4}$$

in the k^{th} profile along the sequence. Suppose further that at every information set *except* the one following an ad "a" posted in state a, the doctor plays her off-equilibrium action with probability ε_k in the k^{th} profile along the sequence. Finally, following an ad "a" posted in

[21] Since it is obviously irrelevant how the patient's strategy is constructed along the sequence, we omit any further reference to it.

state a, the doctor chooses recommendation "β" with probability

$$x_k = \frac{1}{4}\left(\frac{\pi_b\left|\Delta_b^v\right| + \pi_B\left|\Delta_B^v\right|}{\pi_a \Delta_a^v}\right)\varepsilon_k \tag{22}$$

in the k^{th} profile along the sequence. Note that $x_k \in (0,1)$ for each k from (21), and that this sequence of strategy profiles does indeed converge to the equilibrium profile. Now consider the beliefs that the patient will form in each of the scenarios (a)–(c) described above. In scenario (c) – recommendation "α" following an ad with content "s"≠"a" – the posterior probability attached to event $\{a, b, B\}$ in the k^{th} profile along the sequence is:

$$\begin{aligned} & \frac{\pi_a\left(\frac{\varepsilon_k^2}{4}\right) + \pi_b\left(\frac{\varepsilon_k^2}{4}\right) + \pi_B\left(\frac{\varepsilon_k^2}{4}\right)}{\pi_A\left[\frac{\varepsilon_k(1-\varepsilon_k)}{4}\right] + \pi_a\left(\frac{\varepsilon_k^2}{4}\right) + \pi_b\left(\frac{\varepsilon_k^2}{4}\right) + \pi_B\left(\frac{\varepsilon_k^2}{4}\right)} \\ = & \quad \frac{(\pi_a + \pi_b + \pi_B)\,\varepsilon_k}{\pi_A\,(1-\varepsilon_k) + (\pi_a + \pi_b + \pi_B)\,\varepsilon_k} \end{aligned}$$

which converges to 0 as $k \to \infty$. Hence, the limit beliefs make α optimal. In scenario (b) – recommendation "β" following an ad with content "s"≠"a" – the posterior probability attached to state A in the k^{th} profile along the sequence is:

$$\frac{\pi_A\left(\frac{\varepsilon_k^2}{4}\right)}{\pi_A\left(\frac{\varepsilon_k^2}{4}\right) + (\pi_a + \pi_b + \pi_B)\left[\frac{\varepsilon_k(1-\varepsilon_k)}{4}\right]} = \frac{\pi_A\varepsilon_k}{\pi_A\varepsilon_k + (\pi_a + \pi_b + \pi_B)\,(1-\varepsilon_k)}$$

which converges to 0 as $k \to \infty$. Hence, the limit beliefs make β optimal. Finally, consider scenario (a) – recommendation "β" following an ad with content "a". Define

$$\Gamma_k = \pi_A\left(\frac{\varepsilon_k^2}{4}\right) + \pi_a\,(1-\varepsilon_k)\,x_k + (\pi_b + \pi_B)\left[\frac{\varepsilon_k\,(1-\varepsilon_k)}{4}\right],$$

which is the probability of encountering this scenario in the k^{th} profile along the sequence. Then the posterior probabilities attached to states

A, a, b and B in the k^{th} profile along the sequence are, respectively:

$$\frac{\pi_A\left(\frac{\varepsilon_k^2}{4}\right)}{\Gamma_k} = \frac{\pi_A\varepsilon_k}{\pi_A\varepsilon_k + 4\pi_a\left[\frac{(1-\varepsilon_k)x_k}{\varepsilon_k}\right] + (\pi_b+\pi_B)(1-\varepsilon_k)} \tag{A_k}$$

$$\frac{\pi_a(1-\varepsilon_k)x_k}{\Gamma_k} \tag{a_k}$$

$$\frac{\pi_b\left[\frac{\varepsilon_k(1-\varepsilon_k)}{4}\right]}{\Gamma_k} \tag{b_k}$$

$$\frac{\pi_B\left[\frac{\varepsilon_k(1-\varepsilon_k)}{4}\right]}{\Gamma_k} \tag{B_k}$$

Letting $k \to \infty$ and using (22), we observe that

$$(A_k) \to 0$$

Moreover,

$$\begin{aligned} & (a_k)\,\Delta_a^v - (b_k)\,|\Delta_b^v| - (B_k)\,|\Delta_B^v| \\ = \quad & \frac{\varepsilon_k(1-\varepsilon_k)}{4\Gamma_k}\left[\pi_a\left(\frac{4x_k}{\varepsilon_k}\right)\Delta_a^v - \pi_b\,|\Delta_b^v| - \pi_B\,|\Delta_B^v|\right] \\ = \quad & 0 \end{aligned}$$

for each k by virtue of (22). So the limit beliefs make the patient indifferent between α and β.

Summarising the argument so far, we have shown that there is a sequential equilibrium with the strategy profile in Theorem 2 iff $\theta \in [0,1]$ satisfies (17) and (20). Notice that (17) implies

$$\frac{\theta}{1-\theta} \geq \frac{1-\rho(\min\{\gamma, z(0)\}; 0)}{\rho(\min\{\gamma, z(0)\}; 0)}.$$

Moreover,

$$\frac{\theta}{1-\theta} \to \infty$$

as $\theta \to 1$. Since $|\Delta_a^u| < |\Delta_b^u|$, the existence of $\theta \in [0,1]$ satisfying (17) and (20) is therefore equivalent to condition (7).

It remains only to verify that any such SE is Pareto-undominated. Since any such SE delivers patient-optimal treatment in every state, no SE can make the patient better off and any Pareto dominating SE must deliver the same treatment in each state. Since there is no doctor-patient conflict in the current equilibrium, the doctor cannot be better off in any state in any Pareto dominating equilibrium. Thus, if a Pareto dominating SE exists, it must deliver the same treatment in each state and must do so without conflict. It follows easily that in any Pareto dominating SE, firm $\mathcal{A}$ does not advertise in state A, b or B in equilibrium, and firm A's expected advertising expenditure is less than K in state a. It is obvious that we cannot support an equilibrium in which treatment is patient-optimal and firm $\mathcal{A}$ does *no* advertising in equilibrium. Therefore, in any Pareto dominating equilibrium, firm $\mathcal{A}$ must randomise in state a, choosing not to advertise with some probability strictly between 0 and 1. To preserve patient-optimal treatment, the patient response in state a must be the same whether firm $\mathcal{A}$ posts an on-equilibrium ad or does no advertising at all. But then it is strictly optimal for firm $\mathcal{A}$ *never* to advertise in a. This is a contradiction. Therefore, no Pareto dominating SE exists.

This completes the proof of Theorem 2.

C Proof of Theorem 3

A straightforward adaptation of the proof of Theorem 2 suffices, so we just sketch the necessary changes. As before, we start by taking $\theta \in [0,1]$ as given and determine conditions on θ for the indicated strategy profile to constitute (part of) a sequential equilibrium.

Drug firm. The additional condition (8) is required to ensure that the firm has no incentive to withdraw its ad in state b. It plays a role analogous to the one played by Assumption 2 with respect to state a. It is clearly rational for firm $\mathcal{A}$ to advertise truthfully in state a. Condition (8) also implies that $\min\{\gamma, z(0)\} = \gamma$, so firm $\mathcal{A}$'s strategy is sequentially rational iff (8) holds and

$$\theta \;\geq\; 1 - \rho(\gamma; 0) \tag{23}$$

Doctor. To ensure that it is optimal for the doctor to recommend β in state B when firm $\mathcal{A}$ posts an ad with content "a", we need

$$u_B(\alpha) \leq \theta u_B(\beta) + (1-\theta)\left[u_B(\alpha) - c\right] \quad \Leftrightarrow \quad \frac{c}{|\Delta_B^u|} \leq \frac{\theta}{1-\theta} \tag{24}$$

In state $s \in \{a, b\}$, if firm $\mathcal{A}$ posts an ad with content "a", it is optimal to recommend α iff

$$u_s(\alpha) \geq \theta u_s(\beta) + (1-\theta)\left[u_s(\alpha) - c\right] \quad \Leftrightarrow \quad \frac{c}{|\Delta_s^u|} \geq \frac{\theta}{1-\theta} \tag{25}$$

Since $|\Delta_a^u| < |\Delta_b^u|$, conditions (24) and (25) are satisfied iff

$$\frac{c}{|\Delta_B^u|} \leq \frac{\theta}{1-\theta} \leq \frac{c}{|\Delta_b^u|} \tag{26}$$

Patient. The patient's behaviour is optimal in the absence of a detectable deviation from equilibrium provided condition (10) holds. If it did not, then the patient should choose β when observing DTCA with content "a" together with recommendation "α" from the doctor. The detectable off-equilibrium scenarios may be divided into cases (a)–(c) as before. Once again, each scenario can be rationalised by supposing that only one player has deviated. One such rationalisation is possible in scenarios (b) and (c), and two for scenario (a). Suppose, then, that the patient interprets (c) to imply that only firm $\mathcal{A}$ has deviated and hence that the state is A; (b) to imply that only firm $\mathcal{A}$ has deviated and hence that the state is a, b or B; and (a) to imply that either the doctor deviated and we are in state a or b, or that the drug firm deviated and we are in state B. In the latter case, we further suppose that the patient's updated beliefs make him indifferent between treatments. With such beliefs, the patient's strategy is sequentially rational. Furthermore, such beliefs may be shown to be *consistent*. The only change to the construction used in Theorem 2 is the following. If

$$\sum_{s \in \{a,b\}} \pi_s \Delta_s^v \; > \; 0,$$

the sequence $\{\varepsilon_k\}_{k=1}^{\infty}$ satisfies

$$\left(\frac{\pi_B \left| \Delta_B^v \right|}{\pi_a \Delta_a^v + \pi_b \Delta_b^v} \right) \varepsilon_k \; < \; 1$$

for each k and following an ad "a" posted in state a or b, the doctor chooses recommendation "β" with probability

$$x_k \; = \; \frac{1}{4} \left(\frac{\pi_B \left| \Delta_B^v \right|}{\pi_a \Delta_a^v + \pi_b \Delta_b^v} \right) \varepsilon_k$$

in the k^{th} profile along the sequence. If

$$\sum_{s \in \{a,b\}} \pi_s \Delta_s^v \; = \; 0,$$

then $x_k = \sqrt{\varepsilon_k}$ does the needful, as the limit beliefs, conditional on ad "a" and recommendation "β", are just the prior beliefs conditioned on $\{a, b\}$.

Thus, there is a sequential equilibrium with the strategy profile in Theorem 3 iff conditions (8) and (10) hold, and $\theta \in [0,1]$ satisfies (23) and (26). Since $|\Delta_b^u| < |\Delta_B^u|$, the existence of $\theta \in [0,1]$ satisfying the latter two conditions is equivalent to (9).

This completes the proof of Theorem 3.

D Proof of Theorem 4

Suppose the contrary. Thus, if the consumer observes an ad "a", he excludes the possibility that $s \in \{A, B\}$. In each $s \in \{a, b\}$, the consumer must choose α with positive probability, otherwise Firm $\mathcal{A}$ would withdraw its ad (Assumption 4). It follows, given (11), that the doctor must randomise in at least one $s \in \{a, b\}$. Suppose she randomises in both. Since the patient's response to any given recommendation by the doctor must be the same in each state, the doctor can only be indifferent about her recommendation in *both* states if the patient response is independent of the recommendation (Assumption 1). But if the patient's optimal response is the same no matter what recommendation she hears, that response must be to choose β. Given (11), it is not possible for the doctor's strategy to make α an optimal response to *both* recommendations.[22] We have therefore established that the doctor randomises in exactly one $s \in \{a, b\}$. From (11), this state must be b – if the doctor randomised in state a but not state b, the state b recommendation would raise the posterior probability on b above

$$\frac{\pi_b}{\pi_a + \pi_b}$$

and hence the patient would choose β with probability 1 in state b. Thus, the doctor randomises in state b but not in state a. It follows that the recommendation that is made with positive probability in state b but never in state a must elicit drug choice β for certain. If so, the doctor would strictly prefer to make this recommendation in states a and b (Assumption 1). This provides the required contradiction.

[22] If $\theta_s \in (0, 1)$ is the probability that the doctor recommends α in state $s \in \{a, b\}$, then we need $\theta_a > \theta_b$ for α to be an optimal response to "α", and $\theta_a < \theta_b$ for α to be an optimal response to "β".

E Proof of Theorem 5

Since the patient's responses are the same as for Theorems 2 and 3, the proofs of those Theorems provide the conditions for the sequential rationality of the drug firm and physician. In particular, Theorem 2 with $\sigma = 1$ (and hence $\gamma = K + F$) gives the sequential rationality conditions when DTCA is "regulated"; while Theorem 3 with $\sigma = 0$ (and hence $\gamma = K$) gives the sequential rationality conditions when DTCA is "unregulated".

For the drug firm, in addition to Assumptions 4 and 5, sequential rationality requires

$$\theta \;\geq\; 1 - \rho\,(K;0)\,.$$

For the doctor, in addition to Assumptions 1 and 2, we need (20) and (26) to hold, which means

$$\frac{\theta}{1-\theta} = \frac{c}{|\Delta_b^u|} \quad \Leftrightarrow \quad \theta = \frac{c}{c + |\Delta_b^u|}.$$

Next, consider the patient. His on-equilibrium responses are sequentially rational iff condition (12) holds – this ensures α is an optimal choice upon seeing an ad "a" and receiving recommendation "α". We can support his off-equilibrium responses using consistent beliefs constructed as a mixture of those used in the proofs of Theorems 2 and 3. That is, assume the drug firm and doctor "tremble" as per the Theorem 2 construction when Nature chooses DTCA to be "regulated", and as per the Theorem 3 construction when Nature chooses "unregulated" DTCA.

Thus, in summary, besides Assumptions 1, 2, 4 and 5, we require

$$\theta \;=\; \frac{c}{c + |\Delta_b^u|} \;\geq\; 1 - \rho\,(K;0) \quad \Leftrightarrow \quad \frac{c}{|\Delta_b^u|} \;\geq\; \frac{1 - \rho\,(K;0)}{\rho\,(K;0)}$$

to support the SE. This completes the proof of Theorem 5.

References

Beales, H., R. Craswell and S. C. Salop (1981) "The Efficient Regulation of Consumer Information", *Journal of Law and Economics* **24(3)**: 491–539.

Calcott, P. (1999) "Demand Inducement as Cheap Talk", *Health Economics* **8**: 721–733.

Cho, I.K. and D. Kreps (1987) "Signaling Games and Stable Equilibria", *Quarterly Journal of Economics* **102**: 179–221.

De Jaegher, K. and M. Jegers (2001) "The Physician-Patient Relationship as a Game of Strategic Information Transmission", *Health Economics* **10**: 651–668.

Donohue, J.M., M. Cevasco and M.B. Rosenthal (2007) "A Decade of Direct-to-Consumer Advertising of Prescription Drugs", *New England Journal of Medicine* **357(7)**: 673–681.

Frank, R.G., J. Glazer and T.G. McGuire (2000) "Measuring Adverse Selection in Managed Health Care", *Journal of Health Economics* **19(6)**: 829–854.

Keeler, E.B., G. Carter and J.B. Newhouse (1998) "A Model of the Impact of Reimbursement Schemes on Health Plan Choice", *Journal of Health Economics* **17(3)**: 297–320.

Kihlstrom, R.E. and M.H. Riordan (1984) "Advertising as a Signal", *Journal of Political Economy* **92**: 427–450.

Iizuka, T. and G.Z. Jin (2005) "The Effect of Prescription Drug Advertising on Doctor Visits", *Journal of Economics and Management Strategy* **14(3)**: 701–727.

Milgrom, P. and J. Roberts (1986) "Price and Advertising Signals of Product Quality", *Journal of Political Economy* **94(4)**: 796–821.

Narayanan, S., R. Desiraju and P.K. Chintagunta (2004) "Return on Investment Implications for Pharmaceutical Promotional Expenditures: The Role of Marketing-Mix Interations", *Journal of Marketing* **68**: 90–105.

Nelson, P. (1974) "Advertising as Information", *Journal of Political Economy* **81**: 729–754.

Peltzman, S. (1981) "The Effect of FTC Advertising Regulation", *Journal of Law and Economics* **24(3)**: 403–448.

Pitchik, C. and A. Schotter (1987) "Honesty in a Model of Strategic Information Transmission", *American Economic Review* **77(5)**: 1032–1036.

Pitofsky, R. (1977) "Beyond Nader: Consumer Protection and the Regulation of Advertising", *Harvard Law Review* **90(4)**: 661–701.

Rosenthal, M.B., E.R. Berndt, J.M. Donohue, A.M. Epstein and R.G. Frank (2003) "Demand Effects of Recent Changes in Prescription Drug Promotion", Kaiser Family Foundation Working Paper.

Sauer, R.D. and K.B. Leffler (1990) "Did the Federal Trade Commission's Advertising Substantiation Program Promote More Credible Advertising?", *American Economic Review* **80(1)**: 191–203.

Vaithianathan, R. (2006) "Better the Devil You Know than the Doctor You Don't: Is Advertising Drugs to Doctors More Harmful than Advertising to Patients?", *Journal of Health Services Research and Policy* **11(4)**: 235–239.

Wilkes, M.S., R.A. Bell and R.L. Kravitz (2000) "Direct-to-Consumer Prescription Drug Advertising: Trends, Impact, and Implications", *Health Affairs* **19(2)**: 110–128.

Wosinska, M. (2002) “Just What the Patient Ordered? Direct-to-Consumer Advertising and the Demand for Pharmaceutical Products”, Harvard Business School Marketing Research Paper No. 02-04.