Transparency and competition for influence*

Sofia Amaral-Garcia[†], Giacomo Calzolari[‡], Vincenzo Denicolò[§], and Mattia Nardotto[¶]

February 23, 2026

Abstract

We study the impact of mandatory disclosure of contributions paid by interested third parties to decision-makers such as doctors, politicians, or financial advisors. While mandatory disclosure is commonly viewed as a means of reducing potential conflicts of interest, our analysis reveals less benign outcomes when multiple third parties attempt to influence decision-makers in opposing directions. In this case, transparency enables competing third parties to establish separate spheres of influence, where the opposing efforts of rivals do not attenuate their influence. Consequently, decision makers' choices become more polarized. We apply this theory to the market for anticoagulants, using data on doctors' prescriptions, and payments made by pharmaceutical companies to doctors in the United States before and after the Physician Sunshine Act mandated payment disclosure.

Keywords: Lobbying; Transparency; Sunshine Act; Prevalent Influence; Polarization; Drugs; Doctors; Pharmaceutical companies.

J.E.L. numbers: D42, D82, L42

*We thank Alessandro Lizzeri, Matthew Mitchell, Patrick Rey, David Levine, Tom Crossley, Fuhai Hong and Russel Cooper. We are grateful to seminar audiences at UPF (Barcelona), the UNIBG IO Workshop, the EIEF-Bocconi-Unibo workshop, ULB, EUI, Joint Research Center Seville, Bologna, the 2024 Health CEPR 2024, the 2024 Lingnan Conference on Competition (Hong Kong), and the 2024 Oligo Society conference for useful comments. We are, however, fully responsible for any remaining errors.

[†]Université Libre de Bruxelles and i3health, email: sofia.amaralgarcia@ulb.be

[‡]European University Institute and CEPR, email: giacomo.calzolari@eui.eu

[§]University of Bologna and CEPR, email: vincenzo.denicolo@unibo.it

[¶]Université Libre de Bruxelles and CEPR, email: mattia.nardotto@ulb.be

1 Introduction

This paper studies delegated decision-makers (*agents*), who ought to act in the best interests of their stakeholders but receive monetary or in-kind contributions from third parties (*principals*) seeking to influence their decisions. Examples include politicians receiving campaign donations from special interest groups, doctors receiving gifts from pharmaceutical companies, and financial consultants earning fees from issuers of financial assets. (In these cases, the stakeholders are, respectively, voters, patients, and investors.) Such contributions are legal in many countries, but regulations often require their disclosure.[1]

We specifically analyze—both theoretically and empirically—the effects of these transparency regulations. Commonly known as "sunshine laws", such rules are often viewed as a means to promote accountability and reduce the influence of third parties. It is maintained that when contributions are made public, stakeholders become more cautious of agents who receive compensation and more trusting of those who do not. As a result, decision-makers concerned about their reputations may become less willing to accept such contributions, thereby limiting the influence that third parties can exert.

This paper contends that transparency may also have subtler and perhaps unintended consequences that are not as benign as those highlighted by this *reputational* theory. These effects arise when conflicting principals compete to influence an agent in opposing directions. We demonstrate that in such cases, transparency leads agents already inclined toward a specific principal to align even more closely with it, resulting in more extreme choices. In other words, *transparency gives rise to polarization.*

What we exactly mean by polarization may be clarified by considering the market for Novel Oral Anticoagulants (NOACs), which is the focus of our empirical analysis. In this market, multiple pharmaceutical companies supply their own NOAC brand and compete by offering incentives to cardiologists to prescribe their product over those of their competitors. Suppose there are only two branded drugs, and that, absent mandatory disclosure, doctor A prescribes drug 1 to 60% of his patients and drug 2 to the remaining 40%, while doctor B does the opposite. Our theory implies that with mandatory disclosure, doctor A might increase prescriptions of drug 1 to 70% of his patients, while doctor B might do the same with drug 2. In other words, individual doctors' prescribing patterns become more concentrated on a single brand. In politics, polarization would mean that legislators pursue more extreme policies less aligned with the median voter's preferences.[2]

Before explaining the mechanism that produces such polarization, it may be useful to

[1] Regulations mandating the disclosure of financial contributions to politicians date back to the 1930s in the US and have since been adopted in other countries, with increasing stringency. In the pharmaceutical industry, the US Physician Payments Sunshine Act, passed in 2010, requires disclosure of any payment made by pharmaceutical companies to doctors in excess of $10. Similar regulations have been adopted or proposed in several European countries over the past decade. In finance, various provisions, such as the Sarbanes-Oxley Act of 2002, require the disclosure of potential conflicts of interest that can lead to biased decisions or reporting.

[2] Starting from Poole and Rosenthal 1984, this phenomenon – which does not seem to reflect a polarization of the electorate (Fiorina et al. 2008) – has received significant attention in the political science literature. For a comprehensive survey, see Barber et al. (2015).

present some preliminary evidence of its existence. As mentioned, our empirical analysis focuses on cardiologists and the financial contributions they received from manufacturers of different NOAC brands. For each U.S. cardiologist, we observe the prescriptions for each branded drug both before and after the implementation, in 2014, of the Physician Sunshine Act, which made disclosure mandatory at the national level. Earlier disclosure regulations adopted in certain U.S. states serve as a control group where no change in transparency occurred. This setting enables us to conduct a difference-in-differences (DiD) analysis of the concentration of prescriptions at the individual doctor level, measured using the Herfindahl-Hirschman Index (HHI).[3] We perform a similar analysis for the payments doctors receive from pharmaceutical companies, as our theory implies that the same pattern of polarization applies to these contributions. We assess the differences between 2014 (the last year of the no-transparency regime, as the first payment data were publicly released only at the end of the year) and 2018 (the last year considered for this study, for reasons that will be explained below), comparing treated and control states.

Table 1: Difference in differences

	HHI prescriptions			**HHI payments**		
Year	Treated states	Control states	DD	Treatment states	Control states	DD
2014	5217.10	5773.54		5207.85	4509.89	
2018	5450.15	5393.24		5619.82	4281.92	
Diff.	233.05	-380.30	613.35	411.97	-227.97	639.94

Table 1 presents a plain DiD calculation (a more refined analysis with similar findings will be presented later). The results indicate that transparency increased the HHI of prescriptions by more than 10%, and the HHI of payments by almost 15%. The change in the concentration of prescriptions is similar to the numerical example above, where the HHI increases from 5,200 to 5,800.[4]

What is the mechanism that drives polarization? We contend that transparency alters the competition for influence because, under mandatory disclosure, the payments made by a principal can be observed not only by the agent's stakeholders but also by other competing principals. Consequently, a principal may choose to condition its payments to the agent not only on the agent's decisions but also on the level of compensation the agent receives from its rivals. In other words, transparency changes the set of feasible strategies in the influence game.

The primary theoretical contribution of the paper is the characterization of the equilibria

[3] Denoting by s_{id} the share of drug i in doctor d's prescriptions, the HHI of prescriptions for doctor d is $H_d = \sum_{i=1}^{n} s_{id}^2$, where n is the number of drugs. To obtain the overall index of concentration one then takes the average of H_d across doctors.

[4] Guo et al. (2021) document a similar polarization effect of transparency in the market for anti-diabetics. However, their analysis focuses on payments only.

that emerge from these novel strategies. We characterize equilibria where principals offer higher payments on the condition that the agent does not accept significant compensation from rival principals. These strategies split the market into separate spheres of influence, with agents falling into the sphere of influence of a given principal receiving the majority of their contributions from that principal, while other principals are marginalized. It is transparency that enables such quasi-exclusive agreements, as without it, principals would be unable to prevent their competitors from covertly influencing the agent in opposing directions.

Within each sphere of influence, competition among the principals is subdued, if not eliminated. However, principals now seek to expand their respective spheres of influence. This competition *for the market* leads to a distinctive outcome: in equilibrium, each principal attracts those agents whose intrinsic preferences already better align with that principal. Consequently, agents who already lean toward a specific principal tend to further align with it. This produces the polarization effect.

Polarization arises when the interests of principals conflict. However, agents may also make decisions in areas where the interests of principals align. For example, doctors may choose not to treat borderline medical conditions and may opt for older, out-of-patent drugs that are not profitable for the pharmaceutical industry. (For instance, *Coumadin* remains a viable alternative to NOACs.) While each pharmaceutical company promotes its own branded drug, all aim to convince doctors to treat more patients and use branded drugs rather than out-of-patent ones.

Our theory suggests that in areas where the principals' interests align, transparency reduces influence. This is because, within its sphere of influence, each principal limits the rivals' contributions (which are observable), regardless of their objectives (which are not). Consequently, marginalized principals reduce their influence even in areas where all principals share a common goal.

This means that transparency increases distortions in certain choices but mitigates distortions in others. Consequently, the overall welfare effects of transparency in our model are uncertain. To assess the sign and magnitude of these effects, we proceed with a structural estimation of the model for the NOAC market. The estimation reveals, first of all, that cardiologists are quite resistant to influence. For example, consider a doctor with 40 eligible patients who, in the absence of external influences, would treat 50% of them with branded drugs (with the rest being untreated or treated with *Coumadin*), and among those, 50% with a specific brand. To persuade this doctor to prescribe that particular drug to all of his patients, the producer of that brand would have to pay more than $50,000 per year (assuming no countervailing contributions from competing drug suppliers).

Doctors' resistance explains why pharmaceutical companies engage in influencing on a relatively limited scale: on average, each cardiologist receives $1,500 per year from NOACs suppliers. Nevertheless, we find that these contributions significantly impact doctors' choices at the margin. This is because, when a doctor is almost indifferent among different options, even a relatively small inducement can change his behavior.[5]

[5]Several papers have shown that doctors are responsive to the contributions they receive from pharma-

Secondly, the structural estimation shows that the overall welfare effects of transparency are negative. Specifically, the social cost of transparency amounts to [specific amount] per year in the NOACs market. In comparison, an outright prohibition of payments to doctors would increase welfare by [specific amount] per year.

Although the mechanism analyzed in this paper has not, to the best of our knowledge, been previously examined, earlier research has highlighted other potential adverse effects of transparency. For instance, the reputational theory suggests that while transparency may decrease the number of agents receiving payments, those who continue to receive them may obtain higher amounts. The empirical findings of Chen et al. (2019) provide evidence consistent with this prediction.

Inderst and Ottaviani (2009) examine the market for financial advisers and show that mandatory disclosure reduces the overall level of influence exerted by issuers of financial assets. However, it disproportionately reduces the influence of the more efficient issuer, thereby undermining efficiency.

The setting they consider differs from ours, where the agents are the final decision-makers, in that agents (financial advisers) make recommendations to stakeholders (investors), who retain ultimate control over their decisions. In this framework, transparency may have detrimental effects by reducing the informativeness of the recommendations. A similar effect may arise when agents are influencers and stakeholders are their followers, as in Mitchell (2021) and Fainmesser and Galeotti (2021). In these models, influencers may offer both genuine and sponsored recommendations, with transparency regulations mandating the disclosure of sponsored content. Mitchell (2021) and Fainmesser and Galeotti (2021) argue that transparency may actually incentivize influencers to produce more sponsored recommendations. Ershov and Mitchell (2022) bring this theory to the data and find evidence consistent with it.[6]

The remainder of the paper is organized as follows. The next section introduces the model's assumptions. Sections 3 and 4 characterize the equilibria without and with mandatory disclosure of payments, respectively. The theoretical analysis yields several empirical implications, which we explore in the second part of the paper. This begins, in section 5, with a detailed overview of the market for blood thinners and the relevant transparency regulations, followed by a description of our data sources and identification strategy. In section 6, we present and discuss the empirical results of the reduced-form and, in Section 7, those of the structural estimation. Section 8 concludes with suggestions for future research. Several appendices provide supplementary material omitted from the main text, while an online mathematical appendix offers a fully detailed formal analysis of the model.

ceutical companies: see, for instance, xxx.

[6] Cain et al. (2005) conduct an experimental study in a similar setting and find that disclosure can increase bias in recommendations by making advisers feel less morally obligated to act in the best interests of their stakeholders.

2 The model

The model builds upon the theory of special interest groups developed by Grossman and Helpman (1994, 2001). This theory conceptualizes the relationship between agents (such as doctors) and principals (such as pharmaceutical companies) as a common agency game. In this game, the agent is supposed to act in the interest of his stakeholders (e.g., patients),[7] while competing principals seek to influence the agent's choices by offering monetary or in-kind contributions contingent upon the agent's actions. The agent chooses the action that maximizes the sum of his intrinsic utility and the total contributions.

The existing literature assumes that a principal's payments are not observed by its competitors. For our purposes, this represents the no-disclosure benchmark. We contrast it with the case where mandatory disclosure of payments enables principals to condition their contributions not only on the agent's actions but also on the payments received from rivals.

2.1 Players and actions

We consider two competing principals that can influence a number of agents. Assuming agents do not interact with each other, we can focus on the relationship between the principals (indexed by $i = 1, 2$) and a single agent (indexed by $i = 0$).[8]

The agent makes two choices, x and y. This bi-dimensional framework allows for the possibility that the interests of the principals may align along one dimension (x) but conflict along the other (y). Specifically, we assume that both principals benefit from an increase in x, whereas principal 1 desires y to be as large as possible and principal 2 has the opposite preference. Both variables are normalized to the unit interval $[0, 1]$.

For example, in the doctor application, x represents the fraction of patients who receive treatment with newer, branded drugs. The remaining share $1 - x$ are either untreated or treated with older, out-of-patent drugs. On the other hand, y represents the market share of branded drug 1, so $1 - y$ is that of drug 2. (In what follows, we will also use the notation $y_1 = y$ and $y_2 = 1 - y$ to denote these market shares. With this notation, the fraction of the doctor's patients treated with the branded drug i is xy_i.)

2.2 Payoffs

The agent's preferred values of x and y are denoted by ξ and θ, respectively. (For example, ξ represents the fraction of eligible patients that the doctor believes should ideally be treated with a branded drug, while θ indicates the fraction of such patients that the doctor believes should ideally receive drug 1.) Therefore, the value of θ determines the extent to which the agent favors principal 1 over principal 2.

Any deviation from these ideal values causes a loss of utility proportional to the squared distance between the actual choices (x, y) and the bliss point (ξ, θ). Therefore, the agent's

[7] We use masculine pronouns for agents, because over 90% of US cardiologists in the period of study were male, and neutral pronouns for principals.

[8] Appendix D extends the model to the case of more than two principals.

intrinsic utility is, in monetary units:

$$U_0 = -\nu (x-\xi)^2 - \mu (y-\theta)^2 , \qquad (1)$$

where the parameters ν and μ represent the monetary equivalent of the subjective cost of departing from the ideal choices. The convexity of the loss function captures the notion that moving further away from the bliss point is increasingly costly to the agent.[9] The agent maximizes the sum of his intrinsic utility and the payments T_i (which we measure in monetary terms even though they may actually be in-kind) received from the principals:

$$\Pi_0 = U_0 + \sum_{i=1}^{2} T_i. \qquad (2)$$

Consider next the principals. We assume that their payoffs, gross of the payment to the agent, are:

$$U_i = \pi_i x y_i \qquad i = 1, 2, \qquad (3)$$

where π_i is a positive parameter. Thus, both principals gain if x increases. On the other hand, principal 1 wants $y = y_1$ to be as large as possible, while principal 2 has opposite preferences. The net payoff of principal i is:

$$\Pi_i = U_i - T_i \qquad i = 1, 2. \qquad (4)$$

In the doctor example, normalizing the number of patients potentially eligible for treatment to one and interpreting the relevant payoffs in per-capita terms, xy_i represents the volume sales of drug i and π_i the price-cost margin, so U_i is the profit of company i, gross of the payment to the doctor. We take the price-cost margins π_i as given, which amounts to treating prices and production costs of the branded drugs as exogenous for the present analysis.

Since the π_is measure principals' gain from influencing the agent, whereas μ and ν measure the agent's resistance to being influenced, it is intuitive that the equilibrium level of influence will be determined by the ratio between these parameters.

2.3 Strategies

The principals try to influence the agent's behavior by offering conditional transfers. When transparency regulations are absent, such transfers T_i can only be conditioned on the agent's action (x, y). As explained in Grossman and Helpman (1994), such conditioning does not need to be explicitly spelled out in formal agreements; instead, it can be based on informal

[9] Intuitively, when doctors deviate from their optimal decision – such as by increasing the value of x beyond ξ (i.e., over-treating patients) or by increasing the value of y beyond θ (i.e., over-prescribing branded drug 1) – they initially alter the treatment of patients for whom the alternatives are nearly equivalent, so the distortion in choices is not very costly. However, as the distortion increases, doctors begin to change the treatment of patients for whom the next-to-best alternatives are increasingly worse, thereby making the distortions more costly.

understandings and implemented through repeated interactions over time. Analytically, the contribution schedules take the form:

$$T_i = \Gamma_i(x, y) \qquad i = 1, 2. \tag{5}$$

Without loss of generality, we restrict the schedules to be non-negative. We also require that the schedules be sufficiently regular to ensure that the agent's maximization problem has a solution, such as by assuming they are upper semi-continuous.

With mandatory disclosure in place, a principal's payment may depend not only on the agent's actions x and y, but also on whether the agent accepts contributions from the competitor, and the size of such rival contributions: $T_i = \Gamma_i(x, y, T_j)$. We assume that principal i can observe the actual payment T_j but not the schedule offered by the rival, Γ_j. As a result, principal i can condition its payment on T_j but not on Γ_j.[10]

In what follows, we will demonstrate that under transparency, there exist equilibria where the contribution schedules comprise two branches: an upper branch $\Gamma_i^H(x, y)$ that applies if the contribution the agent takes from the rival j is below a certain threshold $\bar{T}_j$ set by principal i, and a lower branch $\Gamma_i^L(x, y)$ that applies if the payment exceeds the threshold:

$$T_i = \begin{cases} \Gamma_i^H(x, y) & \text{if } T_j \leq \bar{T}_j \\ \\ \Gamma_i^L(x, y) & \text{if } T_j > \bar{T}_j. \end{cases} \tag{6}$$

When $\bar{T}_j = 0$, principal i pays the agent only if he does not receive any contribution from the other principal. This is the case of exclusive influence. When $\bar{T}_j$ is positive and finite, (6) represents a "quasi exclusive" agreement. When instead $\bar{T}_j = \infty$, principals effectively do not condition their payment on what the agent receives from the rival, even if they could do so.

These simple schedules are easy to comprehend and do not require the agent to calculate a fixed point. It is important to note that we do not restrict principals to schedules like (6) but show that there exist equilibria where the schedules take this form, even if each principal could use a general schedule $T_i = \Gamma_i(x, y, T_j)$.

2.4 Timing

Principals simultaneously and independently submit to the agent their contribution schedules. After observing these schedules, the agent chooses the contributions to accept, given the principals' offers, and the actions (x, y). Finally, payments are made.

2.5 Information

Appendix A analyzes the model's equilibrium under complete information and establishes two main results. First, multiple equilibria exist. Second, transparency is irrelevant in the

[10] See Szentes (2015) for a model where contract Γ_i may depend on Γ_j. In practice, transparency regulations require the disclosure of payment amounts T_j, while contracts Γ_j may be informal and implicit, and therefore not easily observable.

sense that the Pareto-dominant equilibria for the principals yield identical actions and payoffs regardless of whether disclosure is mandatory. Among these Pareto-dominant equilibria, a particularly salient one is the truthful equilibrium, in which the contribution schedules are linear in x given y, and linear in y given x.

However, Appendix B shows that while observed payment schedules in the NOAC market are linear in x given y, they are distinctly non-linear in y. A parsimonious way to explain this pattern is to assume that the value of θ is the agent's private information, while all other parameters of the model are common knowledge.[11] This also implies that transparency is no longer irrelevant.

Specifically, we assume that while the agent knows the true value of θ, principals only know that θ is uniformly distributed with distribution function:[12]

$$F(\theta) = \frac{1}{\Delta}\left(\theta - \Lambda + \frac{\Delta}{2}\right), \tag{7}$$

and density $f(\theta) = F'(\theta) = \frac{1}{\Delta}$. Thus, θ ranges from $\theta_{\min} = \Lambda - \frac{\Delta}{2}$ to $\theta_{\max} = \Lambda + \frac{\Delta}{2}$, with the parameter Δ representing the size of the support of the distribution, and the parameter Λ the center of the distribution. The agent is on average more inclined toward principal 1 if $\Lambda > \frac{1}{2}$ and toward principal 2 if $\Lambda < \frac{1}{2}$.

For simplicity, the presentation focuses on the symmetric case $\pi_1 = \pi_2 (= \pi)$ and $\Lambda = \frac{1}{2}$.[13]

3 Benchmark: No disclosure

To establish a benchmark for comparison, we begin by analyzing the equilibrium in the absence of transparency. This corresponds to the standard scenario in the literature on special interest groups. Our analysis in this section extends existing models with incomplete information—such as Le Breton and Salanie (2003), Martimort and Semenov (2008a) and Martimort and Stole (2015)—to a setting in which the principals' interests are neither fully aligned nor fully conflicting.

In what follows, we denote a generic function $z(\theta)$ that is only valid within specific upper and lower bounds $\underline{z}(\theta) < \bar{z}(\theta)$ using the notation $\mathcal{I}(\underline{z}(\theta), z(\theta), \bar{z}(\theta))$ defined as:

$$\mathcal{I}(\underline{z}(\theta), z(\theta), \bar{z}(\theta)) = \begin{cases} \underline{z}(\theta) & \text{for } \; z(\theta) \leq \underline{z}(\theta) \\ z(\theta) & \text{for } \; \underline{z}(\theta) \leq z(\theta) \leq \bar{z}(\theta) \\ \bar{z}(\theta) & \text{for } \; z(\theta) \geq \bar{z}(\theta). \end{cases} \tag{8}$$

[11]Given the intractability of multi-dimensional screening models and our focus on competition for influence, it is natural to assume that informational incompleteness concerns the variable capturing the agent's inclination toward the competing principals. For a similar approach, see Le Breton and Salanie (2003), Martimort and Semenov (2008a), and Martimort and Stole (2015).

[12]The uniform distribution allows us to obtain explicit solutions that can be used for the structural estimation, but the qualitative results would extend to any distribution function $F(\theta)$ with monotone hazard rates $\frac{F(\theta)}{f(\theta)}$ and $\frac{1-F(\theta)}{f(\theta)}$.

[13]The formulas for the general case are provided in the mathematical Appendix.

3.1 Monopolistic influence

It is convenient to start with the case where only one principal influences the agent. In this scenario, the principal's problem becomes one of monopolistic screening, which can be solved using standard techniques. The solution is non-degenerate when $\pi^2 < 4\mu\nu$, in which case we obtain:

Proposition 1 *Under monopolistic influence by principal* 1*, the equilibrium actions are:*

$$y^{M_1}(\theta) = \mathcal{I}\left(\theta, 2\nu\frac{\pi\xi - \mu(1+\Delta)}{4\mu\nu - \pi^2} + \frac{8\mu\nu}{4\mu\nu - \pi^2}\theta, 1\right) \tag{9}$$

and

$$x^{M_1}(\theta) = \min\left[\xi + \frac{\pi}{2\nu}y^{M_1}(\theta), 1\right]. \tag{10}$$

When instead $\pi^2 \geq 4\mu\nu$, the incentive to influence the agent is so strong that the monopolistic solution reduces to $y^{M_1}(\theta) = 1$ for all types θ.

The solution for the case of monopolistic influence by principal 2, $y^{M_2}(\theta)$, is obtained by replacing y with $1-y$ and θ with $1-\theta$. Furthermore, the upper bound for $y(\theta)$ becomes θ, while the lower bound is 0.

The payment schedule that supports the monopolistic solutions (detailed in Appendix C) increases linearly in x given y and is quadratic in y. The payment increases with π and decreases with μ and ν, indicating that agents receive lower payments when they are more resistant to influence and higher payments when they are more easily influenced. This property, which will also hold when principals compete, plays an important role in the empirical analysis.

Note that when the solution is interior, $y^{M_1}(\theta)$ increases with θ at a rate more than twice that of the agent's ideal choice, $y = \theta$. This means that principal 1 exerts stronger influence on agents whose preferences already favor it. The reason is that, under incomplete information, influencing "more distant" agents (the low-θ types for principal 1) entails an information rent for the "less distant" ones (the high-θ types for principal 1).[14] Thus, the principal limits its influence on the most distant types to reduce these information rents. This explains the lower bound $y = \theta$ on $y^{M_1}(\theta)$: $y = \theta$ is the agent's choice when the principal stops influencing him along the y dimension, which is precisely what principal 1 may do for sufficiently low-θ types. On the other hand, the principal always influences the agent's choice of x, as there is no private information regarding x and hence no information rents in the x dimension.

Obviously, a principal's unopposed influence pushes the agent's actions in the direction preferred by that principal: $y^{M_1}(\theta) \geq \theta$ and $x^{M_1}(\theta) \geq \xi$. In the doctor example, this means that more patients are treated, and the market share of the influencing company increases, relative to what the doctor would have chosen without external influence.

[14] The *distance* between an agent of type θ and the principals is $1-\theta$ for principal 1 and θ for principal 2.

3.2 Common influence

We now turn to the case in which both principals seek to influence the agent. As is common in models of oligopolistic screening, the equilibrium is derived by positing a specific functional form for the equilibrium schedules—namely, that they are polynomial in x and y—and then verifying that these schedules constitute mutual best responses over the entire space of feasible strategies.[15]

We restrict attention to cases where, in equilibrium, both principals attempt to influence the middle type's choice of y, but in opposite directions. Without this assumption, the principals would not engage in effective competition for influence, and the mechanism analyzed in this paper, through which transparency affects the equilibrium. would no longer operate. This assumption requires two distinct sets of conditions. First, it must hold that $y^{M_1}(\theta) > y^{M_2}(\theta)$, meaning that the agent's choice of y is more favorable to principal 1 when influenced solely by principal 1 than when influenced solely by principal 2. In terms of the model's exogenous parameters, this condition is:

$$\pi > 2(\sqrt{\nu(\Delta\mu + \nu\xi^2)} - \xi\nu). \tag{11}$$

If this condition fails, the equilibrium entails each principal influencing the choice of y only for the types closest to them, while the intermediate types are not influenced at all. In effect, the market would be segmented into two local monopolies even without mandatory disclosure.

Second, it must be the case that if each principal ensures participation of its most distant types, guaranteeing them a positive contribution, this suffices to ensure participation of all types. This is not guaranteed because the agent obtains information rents from both principals. While the rent from each principal decreases monotonically with the distance from the agent, the total rent has a U-shaped profile, potentially placing participation of the intermediate types at risk. In such cases, principals would need to coordinate to ensure participation of these types, and their strategic interaction would turn from competitive to cooperative. This second condition also requires that π be sufficiently large relative to μ and ν. However, the exact formula is too cumbersome to be included here: it is detailed in the mathematical Appendix.

Under this assumption, we obtain:

Proposition 2 *Under no disclosure, the following are equilibrium actions of the common influence game:*

$$y^C(\theta) = \mathcal{I}\left(y^{M_2}(\theta), 3\theta - 1, y^{M_1}(\theta)\right) \tag{12}$$

and

$$x^C(\theta) = \min\left[\xi + \frac{\pi}{2\nu}, 1\right]. \tag{13}$$

[15] The polynomial form is inspired by the structure of transfer schedule in the monopoly case. While this method allows us to identify equilibria of the proposed form, it does not guarantee their uniqueness.

Since both principals now influence the agent's choice of x in the same direction, the equilibrium level of x is the same as if a monopolistic principal had a 100% market share. In the y dimension, however, the principals influence the agent in opposite directions. These opposing efforts do not exactly offset each other (otherwise, the agent would set $y = \theta$, as if he were not influenced at all). This is because, as in the monopoly case, each principal reduces payments to more distant types to limit the information rents obtained by closer types. As a result, each agent type is more heavily influenced by the principal who is closer to him. This creates a counter-clockwise rotation of the schedule $y(\theta)$ compared to the case of no influence where $y = \theta$, which can be seen in Figure 1.

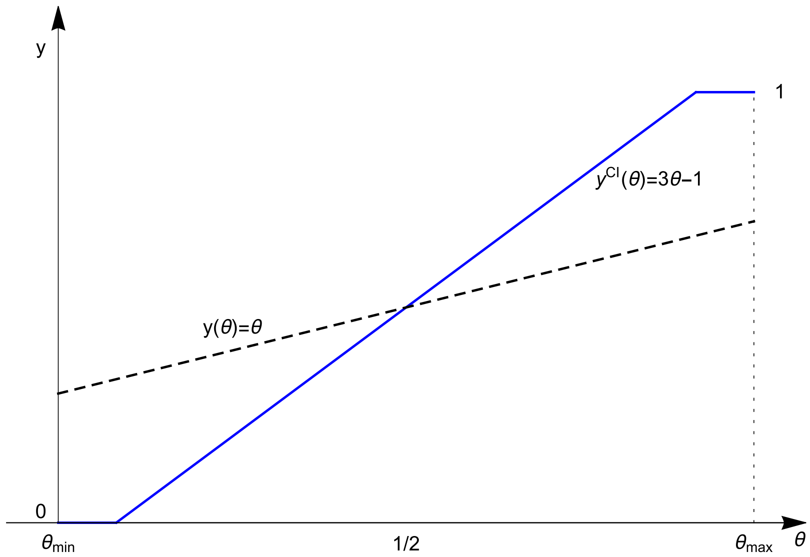

Figure 1: The equilibrium action y as a function of the agent's type θ in the strong competititon case. The blue curve represents the common influence equilibrium $y^C(\theta)$, the black, dashed line the no-influence case $y = \theta$. The figure is drawn for the following parameter values: $\pi = 1$, $\upsilon = 3$, $\mu = 1$, $\xi = 0.8$, and $\Delta = 0.4$.

The exact way in which this effect manifests depends on the intensity of competition for influence. When competition is *strong*, both principals attempt to influence the choice of y of all types. As a result, the agent's choice of y is always subject to countervailing incentives (except for types that are so extreme they select a corner solution where y is either 0 or 1, and are paid only by the principal they favor). This occurs when Δ is small and π is large relative to μ and ν, and it is the case depicted in Figure 1.[16]

The case of *weak* competition arises when the information-rent effect becomes so strong that principals stop influencing the choice of y for the most distant types, and only influence their choice of x, even if y is still strictly between 0 and 1. This occurs specifically when Δ is large and π is small relative to μ and ν. In this scenario, the competition for influence

[16]The exact condition is reported in Appendix C.

is limited to the middle types,[17] while a semi-monopolistic influence regime now prevails at the extremes of the distribution of types. It is *semi*-monopolistic because, although the closer principal influences both the agent's choice of x and y, the more distant principal still influences the choice of x. Thus, extremist types choose the value of y they would if they were influenced only by the closer principal, but select the same value of x as if both influenced them.

The equilibrium payment schedules that support these outcomes, reported in Appendix C, exhibit similar qualitative features to those in the monopolistic influence case. Figure 2 illustrates how the equilibrium contributions depend on the agent's type θ. This dependence results from both the shape of the contribution schedules and the way the agent's equilibrium choices vary with θ. The graph shows that the closer an agent is to a principal, the higher the payment he receives from that principal, with this payment increasing at an accelerating rate. The convexity of payments results from the rising informational rent obtained by closer types. As long as the agent's choice of $y(\theta)$ is interior, each principal continues to pay the agent. However, low-θ agents who set $y = 0$ are paid only by principal 2, and high-θ agents who set $y = 1$ only by principal 1. Naturally, all agents who make these extreme choices receive the same payment, irrespective of their type θ.

4 Mandatory disclosure

We now consider the case where mandatory disclosure is in place, allowing both principals to observe the contributions the agent receives from their competitor. This enables them to condition their payments not only on the agent's actions but also on the rival's contributions.[18]

Since it may not be immediately apparent that principals would seek to take advantage of this possibility, we begin by formally establishing that they actually do.

Proposition 3 *If $\Delta > 0$, the common influence equilibrium described in Proposition 2 is no longer an equilibrium under transparency.*

Intuitively, when the new strategies become available and information is incomplete,[19] if one principal sticks to the equilibrium strategy described in Proposition 2, the rival has a unilateral incentive to deviate. For instance, a principal might introduce an exclusivity

[17] If this effect were even more pronounced, extremist agents (i.e., those with very high or very low θ) would be influenced solely by their nearest principal, while agents in the middle would not be influenced at all. However, condition (11) rules out this scenario, where transparency would have no effect.

[18] To our knowledge, no previous study has investigated these strategies in models of common agency.

[19] The condition $\Delta > 0$ is necessary because, under complete information, the Pareto-dominant truthful equilibrium remains an equilibrium even when the new strategies become available. This follows from the property of *bilateral efficiency*, which ensures that no principal can receive a payoff greater than its marginal contribution to the social surplus—something it can secure without conditioning its payment on the rival's contribution. Otherwise, it would be excluded by a coalition between the agent and the rival principal. However, this property no longer holds under asymmetric information, where informational rents and the resulting distortions undermine bilateral efficiency.

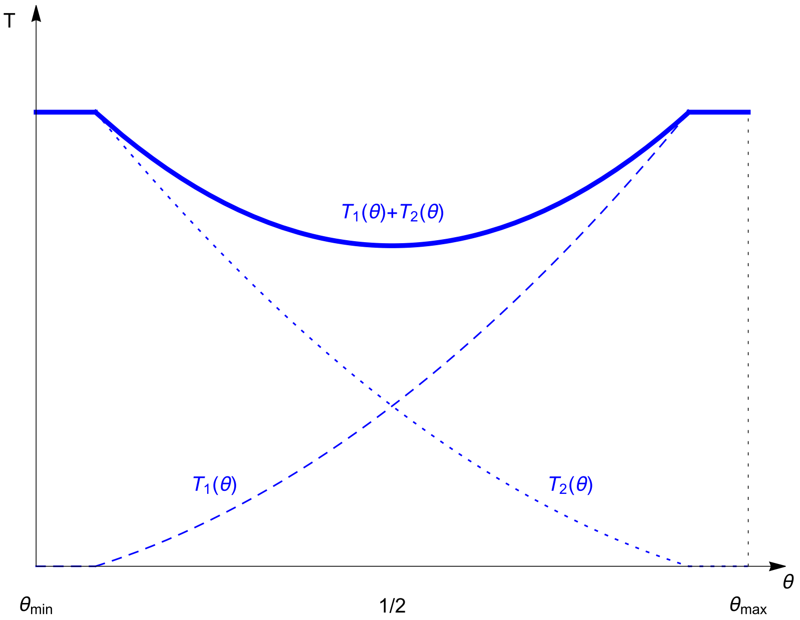

Figure 2: Equilibrium payments under common influence as a function of the agent's type θ. The thick curve represents the total payments $T_1 + T_2$, while the increasing curve represents T_1 and the decreasing curve T_2. The figure is drawn for the same parameter values as Figure 1.

clause—making its payments conditional on the agent not accepting contributions from the rival. Confronted with this condition, the agent must choose between accepting only the deviating principal's payment or only that of the non-deviating one. This has a twofold effect on the deviating principal's payoff. On the one hand, some of the more distant agent types may choose to remain under the rival principal's influence, leading to a loss for the deviating principal, since these types' actions become less aligned with its interests. On the other hand, the closer types will opt into the exclusive relationship, making decisions more favorable to the deviating principal in the absence of opposing influence. The gain from these closer types outweighs the loss from the more distant ones, as each principal primarily values influencing its closest types.

4.1 Exclusive influence

We now proceed to analyze the new equilibria. We begin with the simplest equilibrium arising under mandatory disclosure, in which principals effectively enter into exclusive arrangements akin to the deviation discussed earlier. Although this equilibrium may be Pareto-dominated from the perspective of the principals, it possesses distinctive features that are also present in the more profitable, yet more complex, equilibria we will examine later.

In this exclusive-influence equilibrium, principals offer a positive contribution to the agent only if the agent receives no contribution from the rival. Formally, principals use the following payment schedules:

$$T_i = \begin{cases} \Gamma_i^E(x, y_i) & \text{if } T_j = 0 \\ 0 & \text{if } T_j > 0. \end{cases} \tag{14}$$

Clearly, if principal i offers only an exclusive contribution schedule, principal j has no incentive to offer a non-exclusive schedule, making the exclusive-influence equilibrium self-sustaining.[20] This observation leads directly to the following result:

Proposition 4 *Under transparency, there exists an* ***exclusive-influence*** *equilibrium in which principals offer schedules of the type (14). In this equilibrium, the population of agents is divided into two spheres of exclusive influence: types* $\theta \in \left[\theta_{\min}, \frac{1}{2}\right)$ *accept only the contribution of principal 2, while types* $\theta \in \left(\frac{1}{2}, \theta_{\max}\right]$ *that of principal 1.*[21]

The equilibrium level of y *is:*

$$y^E(\theta) = \begin{cases} y^{M_2}(\theta) & \text{for } \theta < \frac{1}{2} \\ y^{M_1}(\theta) & \text{for } \theta > \frac{1}{2}. \end{cases} \tag{15}$$

and that of x *is:*

$$x^E(\theta) = \begin{cases} x^{M_2}(\theta) & \text{for } \theta < \frac{1}{2} \\ x^{M_1}(\theta) & \text{for } \theta > \frac{1}{2} \end{cases} \tag{16}$$

The exclusive-influence equilibrium creates two distinct spheres of influence, each comprised of agents who are already predisposed toward a particular principal. As a result, principal 1 influences only high-θ types, while principal 2 focuses solely on low-θ types. Within its sphere of influence, each principal acts as a monopolist, leading to an allocation identical to that under monopolistic influence by that principal alone. This equilibrium is thus characterized by both polarized contributions and polarized actions. Contributions are polarized, as each agent type receives a positive payment from only one principal. Actions are likewise polarized: y is higher than under common influence for high θ and lower for low θ. In other words, the schedule $y^E(\theta)$ rotates counterclockwise relative to the equilibrium schedules under no disclosure, $y^C(\theta)$, as illustrated in Figure 3.

[20]Note the distinction between the exclusive-influence strategies described in equation (14) and the exclusive-dealing contracts studied in the Industrial Organization literature. Exclusive-influence schedules promise a positive contribution to agents only if they do not receive any payments from rivals, whereas exclusive-dealing arrangements permit buyers to purchase a firm's product only if they refrain from buying from competitors. In our framework, this would be analogous to conditioning contributions on agents not endorsing rival products, such as doctors not prescribing competing drugs. These strategies would be feasible even with no mandatory disclosure, but they do not constitute an equilibrium in our model.

[21]The middle type $\theta = \frac{1}{2}$ is indifferent between being influenced by principal 1 or 2.

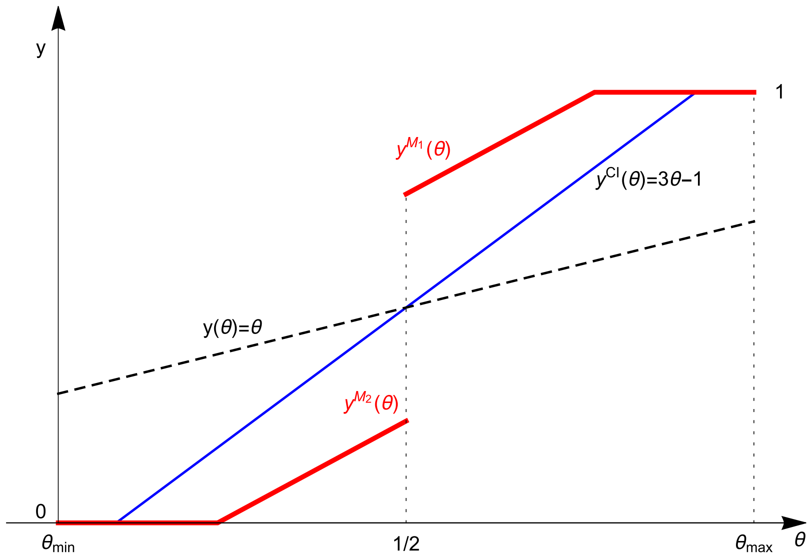

Figure 3: The equilibrium value of y as a function of agent's type θ in the exclusive-influence equilibrium (in red) and in the common-influence equilibrium (in blue). The back dashed line represents the case where the agent is not influenced, $y = \theta$. The figure is drawn for the same parameter values as Figure 1.

Although the exclusive-influence equilibrium may resemble a collusive agreement in which the two principals divide the market into local monopolies, competition between them persists. However, the nature of this competition changes: rather than competing *in the market*, the principals now compete *for the market*—each seeking to expand its sphere of influence by attracting more agents to its side. The competition for the market takes place through the choice of the constant terms of the payment schedules, denoted by α_{0i}.[22] Increasing α_{0i} has two contrasting effects on i's payoff. On one hand, a higher α_{0i} enlarges i's sphere of influence, and a principal's payoff is higher within its sphere of influence than outside it. On the other hand, a higher α_{0i} increases the payments to agents already within i's sphere of exclusive influence. For a given α_{0j}, the optimal level of α_{0i} balances these effects, determining principal i's best response to α_{0j}. The intersection of the best response functions of both principals defines the equilibrium.

Regarding x, exclusive influence results in a lower equilibrium level compared to the common influence scenario. This is because the agent is influenced by only one principal in his choice not only of y, where the principals' interests conflict, but also of x, where they align. Since each principal would influence x proportionally to its market share, the absence of a principal's contribution reduces the overall influence on x. In the doctor example, this

[22] The variable components are not affected by this competition as they are determined by the hazard rates $\frac{F(\theta)}{f(\theta)}$ and $\frac{1-F(\theta)}{f(\theta)}$, which remain the same regardless of the size of the spheres of influence, as long as each principal's sphere of influence includes its closest types.

means that transparency reduces the fraction of patients treated with branded drugs. This effect is particularly pronounced for intermediate θ types, where the market share of the sole influencing principal is smallest. This can be observed in Figure 4.

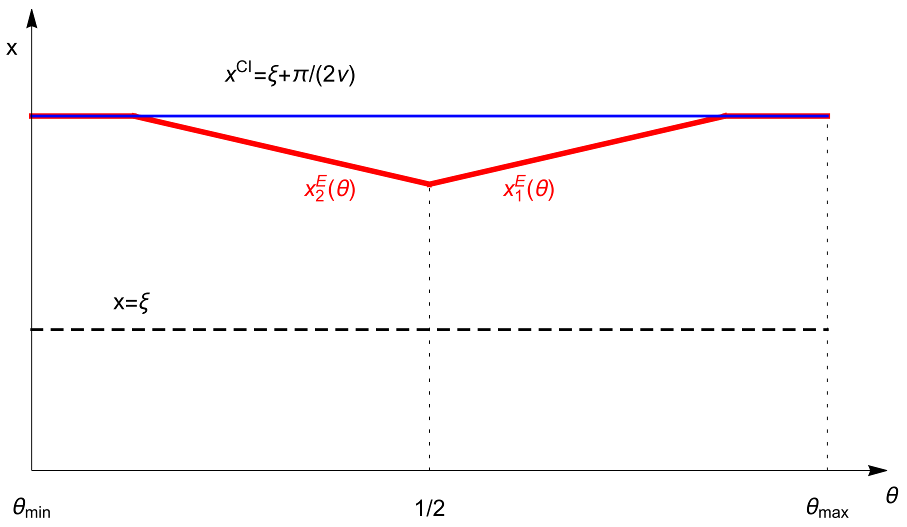

Figure 4: The equilibrium level of x as a function of agent's type θ in the exclusive-influence equilibrium (in red) and in the common-influence equilibrium (in blue). The back, dashed line represents the case of no influence, $x = \xi$. The figure is drawn for the same parameter values as Figure 1.

4.2 Prevalent influence

As noted, the exclusive-influence equilibrium described in Proposition 4 may be Pareto-dominated from the perspective of the principals. This is because each principal prevents its rival from influencing not only y, where their interests conflict, but also x, where their interests align. This suggests that there is room for improvement from the principals' viewpoint.

Improving upon the exclusive-influence equilibrium requires that each principal, within its sphere of influence, allows its rival to incentivize the agent's choice of x while still preventing its influence over y. The challenge is that each principal can condition its payments on the rival's contributions, but not on the motivations behind those contributions. Nevertheless, a limited form of coordination can emerge even in a non-cooperative equilibrium. The following proposition characterizes a class of equilibria in which agents remain divided into distinct spheres of influence, but the nature of this influence shifts from *exclusive* to *prevalent*. Under prevalent influence, agents who fall within principal i's sphere of influence still receive a positive contribution from principal j, but this contribution is subject to a cap. As a result, each agent has both a primary and a secondary contributor. While the cap may vary—giving rise to a continuum of equilibria—it is always sufficiently low to ensure that the secondary contributor affects only the agent's choice of x, and not of y.

Proposition 5 *Under mandatory disclosure, if* $2(1-\Delta)\mu\nu - \nu\xi - \pi^2 > 0$*, there exists a continuum of symmetric* ***prevalent influence*** *equilibria in which principals use contribution schedules of the form (6), with* $\bar{T}_i = \bar{T}_j = \bar{T}$*. The variable* $\bar{T}$ *can take on any value in an interval* $\left[0, \hat{T}\right]$*, where*

$$\hat{T} = \max\left[0, \pi^2 \frac{\left[2(1-\Delta)\mu\nu - \nu\xi - \pi^2\right]^2}{4\nu\left(\pi^2 - 4\mu\nu\right)^2}\right], \tag{17}$$

with equilibria having higher $\bar{T}$ *Pareto dominating, from the principals' viewpoint, those with lower* $\bar{T}$*. There are no equilibria where* $\bar{T} > \hat{T}$*.*

In all of these equilibria, the population of agents is divided into two spheres of prevalent influence: types $\theta \in \left[\theta_{\min}, \frac{1}{2}\right)$ *accept the large contribution* $\Gamma_2^H(x,y)$ *from principal 2 and the small contribution* $\Gamma_1^L(x,y)$ *from principal 1, while types* $\theta \in \left(\frac{1}{2}, \theta_{\max}\right]$ *accept the large contribution* $\Gamma_1^H(x,y)$ *from principal 1 and the small contribution* $\Gamma_2^L(x,y)$ *from principal 2.*

The equilibrium level of y *is the same as in the exclusive-influence equilibrium,* $y^P(\theta) = y^E(\theta)$*, while the level of* x *is higher:*

$$x^P(\theta) = x^E(\theta) + \sqrt{\frac{\bar{T}}{\mu}} \tag{18}$$

The condition $2(1-\Delta)\mu\nu - \nu\xi - \pi^2 > 0$ ensures that the equilibrium value of $y(\theta)$ remains interior for all *theta*. If this condition is not satisfied, the exclusive-influence equilibrium is the only one that exists. Meanwhile, the upper bound $\hat{T}$ on $\bar{T}$ guarantees that $x^P(\theta)$ does not exceed $x^C(\theta)$.

Some features of the prevalent-influence equilibria described in Proposition 5 are already familiar from the earlier analysis. As in the case of exclusive influence, an agent's choice of y is determined solely by the principal within whose sphere of influence the agent falls in, resulting in the same equilibrium level of y as under exclusive influence. Once again, principals compete for expanding their spheres of influence through the fixed components of their contribution schedules.

However, other aspects of these equilibria are novel. Under prevalent influence, each agent receives contributions from both a primary and a secondary principal, with the payment from the latter constrained by a cap imposed by the former. This raises two questions. First, how is the secondary contributor's payment schedule designed? (The primary contributor's schedule is the same as under monopolistic influence, *modulo* a constant.) Second, how is the cap set by the primary contributor?

Regarding the first question, the secondary contributor's schedule, $\Gamma^L(x, y)$, is designed firstly to comply with the cap set by the primary contributor. This means that while principals seek to expand their spheres of influence, they ultimately recognize that only the closest agent types can be attracted. Thus, in equilibrium, they accept the role of secondary contributor for more distant types. Secondly, the secondary contributor's schedule is designed in such a way as to influence only the agent's choice of x, not y.

This latter property is crucial for understanding how the cap $\bar{T}$ is set in equilibrium. The cap limits how much influence the secondary contributor can exert on x: the greater the room allowed, the stronger its influence and the higher its payoff. The primary contributor does not benefit from the increase in x due to the influence exerted by the rival, as in equilibrium its payment increases with x at the same rate as its revenue. However, it is not harmed by that influence either and is, therefore, willing to tolerate it provided the secondary contributor does not attempt to influence the agent's choice of y. This is precisely what the secondary contributor does when $\bar{T} \leq \hat{T}$. If $\bar{T} > \hat{T}$, however, the secondary contributor would try to sway the agent's choice of y, which would harm the primary contributor. Such behavior is not tolerated in equilibrium, and thus equilibria where $\bar{T} > \hat{T}$ do not exist.

As long as $\bar{T} \leq \hat{T}$, the game becomes one of pure coordination with multiple equilibria: given $\bar{T}_i$, principal i is indifferent among different levels of $\bar{T}_j$ and thus, in equilibrium, can choose any value of $\bar{T}_j \leq \hat{T}$, and the same is true of principal j. (While the proposition specifically focuses on symmetric equilibria where the caps $\bar{T}_i$ and $\bar{T}_j$ coincide, there are also asymmetric equilibria where $\bar{T}_i \neq \bar{T}_j$.) Since each principal is both a primary contributor (in its sphere of prevalent influence) and a secondary contributor (in the sphere of prevalent influence of the rival), a coordinated increase in the caps is jointly beneficial. Consequently, the Pareto-dominant equilibrium is the one in which the caps are set at their maximum sustainable level, $\bar{T} = \hat{T}$.

In prevalent-influence equilibria, the polarization effect of transparency on prescriptions mirrors that observed under exclusive influence. The polarization of contributions persists but is attenuated, as agents now receive positive payments from both principals. Transparency still mitigates principals' influence on x, though to a lesser extent than in the case of exclusive influence.

4.3 Welfare effects of transparency

Who gains and who loses from mandatory disclosure? Conventional wisdom holds that transparency undermines the principals' ability to influence the agent, thereby reducing their payoffs while benefiting the agent's stakeholders. However, our theory yields more nuanced—and potentially counterintuitive—conclusions.

We begin by examining the impact of transparency on social surplus S, defined as the sum of the players' payoffs: $S = U_0 + U_1 + U_2 = U_0 + \pi x$.[23]

Corollary 1 *Transparency always decreases the social surplus S.*

[23] Note that the payments made by the principals to the agent cancel out in the social surplus calculation.

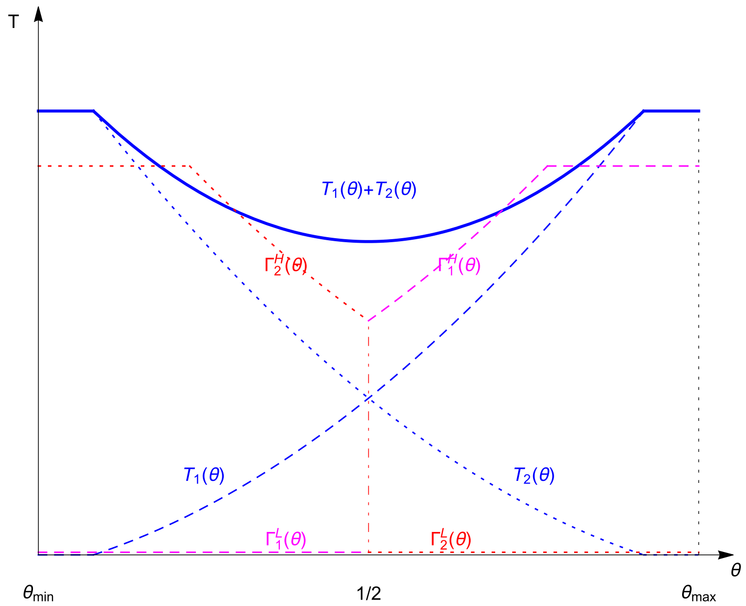

Figure 5: Equilibrium payments as a function of agent's type θ in the prevalent-influence equilibrium (in red or magenta) and in the common-influence equilibrium (in blue). The figure depicts an exclusive prevalent equilibrium where $\bar{T} = 0$ and is drawn for the same parameter values as Figure 1.

The proof of the corollary is simple. Social surplus is maximized when the agent's choices are $y(\theta) = \theta$ and $x(\theta) = \min\left[\xi + \frac{\pi}{2\nu}, 1\right]$. In the common-influence equilibrium, the agent sets x at the efficient level, while y is excessively polarized. Mandatory disclosure distorts both margins: it lowers x below its efficient level and further amplifies the polarization of y. Both distortions reduce the overall surplus.

This result already casts doubt on the view that the welfare effects of mandatory disclosure are necessarily benign. However, one might question whether social surplus S is the appropriate welfare criterion. A plausible alternative is to focus solely on the welfare of stakeholders. This may be captured by the function U_0, which reflects the agent's genuine assessment of how his choices affect stakeholders.

For example, U_0 can be interpreted as a doctor's informed judgment about the appropriateness of available treatments for his patient population. If the doctor is the most qualified expert to make such judgments, then his bliss point (ξ, θ) can be interpreted as the ideal outcome for those patients. However, doctors may adopt an overly conservative stance on the value of branded drugs—either because they are naturally cautious and underestimate the benefits of innovative therapies,[24] or because they internalize the financial cost of prescribing branded drugs without accounting for the fact that pharmaceutical profits may incentivize innovation and thus are not entirely socially wasteful.

A general approach that accommodates these differing perspectives is to define welfare as a weighted sum, $U_0 + \lambda \pi x$, where the parameter λ ranges from 0 to 1. The case $\lambda = 0$ corresponds to a world in which the doctor accurately evaluates the true social costs and benefits of new treatments. At the other extreme, $\lambda = 1$ recovers the definition of social surplus used earlier.

Under this broader notion of social welfare, the effects of transparency become ambiguous when λ is sufficiently low. Even in the extreme case where $\lambda = 0$, however, transparency improves social welfare only if ν is very large and ξ is very small. Intuitively, the benefit of transparency in this case is that it reduces the over-treatment $x - \xi$. However, for this benefit to outweigh the welfare loss from increased polarization in y, the marginal social value of reducing x must be sufficiently high.

The effect of transparency on the principals' payoffs in our model is also somewhat surprising. While conventional wisdom suggests that mandatory disclosure harms the principals by constraining their ability to influence the agent, our model shows that it may in fact benefit them. The key mechanism is that transparency shifts the nature of competition from "in the market" to "for the market"—and the latter may be less intense. This occurs particularly when Δ is large, indicating a high degree of asymmetry in information about the doctor's intrinsic preferences.[25]

Ultimately, therefore, the effect of transparency on both individual payoffs and overall welfare remains theoretically ambiguous. To evaluate the sign and magnitude of these effects in a concrete setting, we structurally estimate the model using data from the NOACs market. Before turning to the estimation, we first provide a more detailed overview of the NOACs market and present reduced-form evidence on the effects of transparency.

5 The market for blood thinners

In this section, we describe our empirical setting: the U.S. market for New Oral Anticoagulants (NOACs), commonly known as blood thinners. This is one of the most lucrative markets in which multiple branded drugs compete directly, with total annual sales of approximately $25 billion, and two products consistently ranking among the top ten best-selling drugs. It is also one of the three largest markets in terms of total payments made by

[24] In health economics, it is indeed not uncommon to estimate the welfare effects of different treatments based on the opinions of super experts rather than that of average practitioners: see, for instance, Allcott and Taubinsky (2015), Bronnenberg et al. (2015), and Handel and Kolstad (2015).

[25] The possibility that principals may benefit from transparency helps explain why some pharmaceutical companies have voluntarily disclosed payment information even in the absence of regulatory mandates. While such behavior is often interpreted as a strategy to preempt potential regulation, our analysis suggests that transparency can, under certain conditions, serve the principals' own interests.

pharmaceutical companies over the period we study.[26]

5.1 Market structure

Blood thinners are medications that prevent or slow the formation of blood clots and are commonly prescribed to treat or prevent conditions such as deep vein thrombosis, stroke, and pulmonary embolism.[27] For decades, the market was dominated by Warfarin (brand name *Coumadin*), which was approved by the FDA in 1954. After Warfarin's patents expired in the 1980s, generic versions entered the market, and the competitive landscape remained stable for several decades. This changed in the early 2010s with the FDA approval of the first New Oral Anticoagulants (NOACs): Dabigatran (*Pradaxa*) in 2010, Rivaroxaban (*Xarelto*) in 2011, and Apixaban (*Eliquis*) in 2012.[28]

NOACs offer a key advantage over Warfarin: they provide more predictable anticoagulant effects, which reduces the need for frequent patient monitoring.[29] Despite their higher cost, this clinical convenience has gradually eroded Warfarin's market share, though it remains in use to this day.

In the mid-2010s, two additional NOACs received FDA approval: Endoxaban (*Lixiana*) in 2015 and Betrixaban (*Savaysa*) in 2017. Although their initial market penetration was modest, their sales began to rise significantly after 2018. Accordingly, our analysis focuses on the period from 2014—when payment data first became available following the implementation of the Sunshine Act of 2010—to 2018, prior to the changes in market structure caused by the entry of these newer drugs.

The list prices of all branded NOACs remained closely aligned throughout the study period and continue to be similar today.[30] In the absence of publicly available data on production costs and discounts from list prices, and lacking strong evidence of systematic differences among the drugs along these dimensions, it is reasonable to assume that the profitability parameter π_i is roughly constant across all branded products.

However, the evolution of market shares, documented in Appendix F, reveals marked asymmetries. One such asymmetry arises from safety concerns associated with *Pradaxa*,

[26] Statins and anti-diabetics represent two other markets of comparable size. However, they involve a significantly larger number of branded competitors, making the application of our model less straightforward. That said, the polarization in payment patterns we observe in the anticoagulant market has also been documented in the anti-diabetic market by Guo et al. (2021), as noted earlier.

[27] Our analysis focuses exclusively on oral anticoagulants—drugs administered in pill or tablet form. For intravenous treatments, *Heparin* remains the dominant option, and no branded alternatives hold significant market share.

[28] In the US, these drugs are suppliedby Boehringer Ingelheim (*Pradaxa*), Janssen (*Xarelto*), and Bristol Myers Squibb (*Eliquis*), respectively.

[29] Monitoring for NOACs typically occurs before treatment begins, again 1–3 months after initiation, and then every 6–12 months. In contrast, Warfarin requires blood testing every 2 to 4 weeks. NOACs also involve fewer dietary restrictions.

[30] The list prices of *Eliquis* and *Xarelto* increased in parallel from approximately $300 per month in 2014 to over $400 per month in 2018, prompting speculation about potential collusive behavior. See the 2022 report by *Patients for Affordable Drugs*, which can be found at https://patientsforaffordabledrugs.org/wp-content/uploads/2022/04/Eliquis-and-Xarelto-report.pdf).

the first NOAC to receive regulatory approval. The drug's anticoagulant effects proved less predictable in certain patient populations than initially claimed, and less predictable than those of *Xarelto* and *Eliquis*. As a result, *Pradaxa* became the subject of successful damage lawsuits brought by patients, and its market share has declined steadily since 2014, being now largely confined to patients who began treatment before the introduction of *Xarelto* and *Eliquis*. This pattern indicates that physician preferences for *Pradaxa* are more stable and predictable than those for the other two drugs. This suggests that the asymmetric three-principal version of the model, presented in Appendix E and yielding the same qualitative conclusions as the baseline specification, may offer a good fit for this market structure.

Another asymmetry concerns the two leading branded drugs. During the study period, *Xarelto* initially led in sales but was eventually surpassed by *Eliquis*, the last of the three NOACs to enter the market (see Appendix F). This pattern suggests that the distribution function $F(\theta)$ may have shifted over time, with the median doctor initially favoring *Xarelto* and gradually moving toward *Eliquis*. This dynamic plays an important role in the structural estimation presented in Section X.

5.2 Data

Our analysis focuses on U.S. cardiologists holding an active medical license. For each physician, we observe all Medicare Part D prescriptions for Warfarin and each of the three branded NOAC, as well as all payments received from pharmaceutical companies.

Data on prescriptions are drawn from the Medicare Part D Prescriber File, published by the Centers for Medicare & Medicaid Services (CMS). This dataset is organized at the provider–drug level and contains detailed information on medications prescribed to Medicare Part D beneficiaries. Each observation includes the prescribing physician's name, the prescribed drug, the total cost of all associated claims, and the number of claims filed.

Data on payments are obtained from the Open Payments database, also maintained by CMS, and established under the Sunshine Act of 2010 (see below). This dataset records transfers of value from pharmaceutical companies to physicians, including the physician's unique identifier, the date of payment, and, in many cases, the name(s) of the drug(s) linked to the payment. Transfers may take various forms, such as cash, meals, travel, speaking fees, honoraria, gifts, or research funding. For non-cash contributions, the database includes their dollar value.

We merge these two datasets using the National Provider Identifier (NPI), a unique ID assigned to each physician. Additional physician-level information—including specialty, gender, geographic location, year of medical school graduation, and whether the physician graduated from a top-20 U.S. medical school (as ranked by *U.S. News & World Report* in the Research category)—is obtained from the Doctors and Clinicians National Downloadable File, formerly known as the *Physician Compare* dataset, also provided by CMS.[31]

[31]We have downloaded several waves of *Physician Compare* to cover the years that we span with our data: December 2014, July 2015, April 2016, and December 2018. For waves that are no longer available on the official webpage of *CMS Data* (see https://data.cms.gov/provider-data/archived-data/doctors-clinicians), we have resorted to the NBER repository (see https://data.nber.org/compare/

.

5.3 Sample

Starting from the universe of 21,684 physicians specializing in cardiology, we apply three selection criteria to ensure data quality. First, we restrict the sample to physicians who prescribed at least $5,000 worth of anticoagulants per year to Medicare Part D beneficiaries (by way of comparison, the average doctor in our final sample prescribes a yearly total of approximately $110,000). Given the cost of annual treatment, this threshold implies that each physician treated at least two patients with NOACs, ensuring the data are meaningful for analyzing polarization. Second, we retain only those physicians who can be matched to the *Physician Compare* dataset. Third, we exclude physicians whose prescribing behavior is unlikely to reflect autonomous decision-making. This includes individuals who had not yet graduated from medical school during the study period (and were thus likely practicing under supervision), as well as those who obtained their degree more than 60 years before the start of the data, and who are likely to delegate prescribing decisions to younger colleagues.

Our final sample comprises 16,674 physicians who meet all three criteria. The vast majority of the 5,010 excluded physicians had little or no NOAC prescribing activity during the study period. Indeed, many of the 236 physicians who could not be matched with *Physician Compare*, and the 438 classified as either "too young" or "too old," would already have been excluded based on the first criterion.[32]

5.4 Descriptive statistics

Appendix G reports detailed descriptive statistics for our sample. Here, we briefly note that only 10% of the doctors in our sample are female, which is why we use masculine pronouns to refer to agents throughout the paper.[33]

As previously noted, the average physician in our sample prescribes approximately $110,000 worth of anticoagulants annually. The vast majority of this spending is accounted for by the three branded NOACs, as Warfarin is substantially less expensive by comparison. Over 80% of cardiologists prescribe three or all four available anticoagulants (i.e., the three branded NOACs and Warfarin).

On average, physicians receive nearly $1,000 in payments from pharmaceutical companies. However, 37% of doctors receive no payments at all. Among those who do receive contributions, the average payment is approximately $1,500. Nearly half of these recipients receive payments from all three pharmaceutical companies, while only about one-quarter receive payments from a single firm.

`physician/2016/12/physician-compare`).

[32] For example, the 236 doctors whom we cannot match prescribe on average $6,644 per year.

[33] The gender imbalance is particularly pronounced in cardiology relative to other specialties. However, the average year of graduation for female doctors in our sample is 1995, compared to 1988 for the full sample. This suggests that the share of female cardiologists is likely to grow in the future.

6 Reduced-form evidence

This section presents reduced-form evidence that speaks to the empirical implications of our theory. The following section turns to the structural estimation of the model.

6.1 Preliminary validation

We begin by documenting a strong, positive correlation between industry payments and physicians' prescription behavior. This serves as a preliminary validation of our model: if company payments were entirely uncorrelated with prescribing patterns, the Grossman and Helpman (1994) model of lobbying would not provide a meaningful description of the market and, therefore, could not serve as a foundation for evaluating the impact of transparency.

It is important to emphasize that our theory does not imply a causal relationship between payments and prescriptions, as both variables are endogenously determined. However, the model does predict a positive correlation between the value of prescriptions for product i made by doctor d in year t, and the amount paid by company i to doctor d.[34]

Table 7 shows that this correlation indeed exists and persists even after controlling for doctor and drug fixed effects to mitigate concerns about spurious correlation. The inclusion of fixed effects reduces the estimated coefficient, suggesting that payments are partly driven by time-invariant drug and physician characteristics. Nevertheless, the coefficient remains both economically meaningful and statistically significant even after accounting for both fixed effects (column 3), so that the remaining variation captures doctor–drug–time–specific changes in payments.[35]

As an additional robustness check, we implement the correction method proposed by Oster (2019), based on Altonji et al. (2005) and Altonji et al. (2010). The resulting adjusted coefficient is 0.78, confirming the presence of a strong positive relationship between payments and prescription behavior.

6.2 Identification strategy

Having established that payments and prescriptions are related in a manner consistent with our theoretical model, we now turn to examine the specific effects of mandatory disclosure. The analysis leverages the first implementation, in 2014, of the Physician Payments Sunshine Act. Passed by the U.S. Congress in 2010,[36] the legislation requires pharmaceutical and medical device manufacturers to report to the CMS any payment or in-kind contribution exceeding \$10 per instance or \$100 annually made to physicians and other health care providers. CMS, in turn, discloses this information to the public. In principle, manufacturers are also required to specify the drug or device associated with each contribution.

[34] We are not the first to document such a correlation. See, for instance, Agha and Zeltzer (2022)) for the market we study, as well as Carey et al. (2021) and Grennan et al. (2021) for different pharmaceutical markets.

[35] For a similar approach, see Agha and Zeltzer (2022).

[36] The Act was bi-partisan. Its stated goal was to increase transparency in the relationship between physicians and pharmaceutical and medical device manufacturers

Table 2: Marginal returns on payments

Dependent variable: Amount prescribed					
		All years		**Pre-trasparency**	**Post-trasparency**
	(1)	(2)	(3)	(4)	(5)
Amount received	2.452***	2.268***	1.721***	1.021***	1.954***
	(0.183)	(0.170)	(0.144)	(0.202)	(0.143)
Year FE	YES	YES	YES	YES	YES
Doc FE	NO	YES	YES	YES	YES
Drug FE	NO	NO	YES	YES	YES
R^2	0.080	0.302	0.488	0.553	0.528
Observations	319528	319528	319528	59660	259868

Notes: *Amount prescribed* is the yearly total amount of NOACC prescribed by the doctor from a single manufacturer; *Amount received* is the yearly total amount of payments received by the doctor from the same manufacturer. Columns (1) to (3) make use of the full sample while columns (4) and (5) focus on the pre-transparency period and the post-transparency period respectively. Standard errors are clustered at the physician level. *, **, and *** denote significance at the 10, 5, and 1 percent level, respectively.

Payment data are released annually. As a result, the strategic behavior analyzed in the first part of the paper—whereby pharmaceutical companies may condition payments to physicians on the payments those physicians receive from rival firms—could only become viable in the year following the release of the first dataset. This feature allows us to observe both prescribing and payment patterns in the final year of the pre-transparency regime, as well as during multiple years of the post-transparency regime. Comparing outcomes across these periods provides a basis for empirically testing the implications of our theoretical model.

Beyond enabling a simple pre-post comparison, our data also supports a difference-in-differences (DID) analysis. This is because at the time the Sunshine Act took effect, six U.S. states already had some form of disclosure regulation in place. However, three of these states (Maine, West Virginia, and the District of Columbia) only required payments to be reported to the state but not made public, which did not effectively create transparency as defined in this paper. In contrast, Massachusetts, Minnesota, and Vermont had regulations that were comparable to, or stricter than, those in the Sunshine Act. Accordingly, we treat these three states as a control group, while the remaining states where payment data were published for the first time in 2014 constitute the treatment group.

A similar identification strategy has been employed in other studies examining the impact of the Sunshine Act.[37] Two important caveats, however, should be noted. First, in our setting, the control group—physicians located in states with pre-existing transparency regulations—does not consist of "untreated" subjects, but rather of physicians who had already been "treated."[38]

[37] Such studies include Guo et al. (2021), Carey et al. (2021) and Agha and Zeltzer (2022).

[38] Although less common, this framework has been employed in situations similar to ours, where laws or

Second, although CMS currently publishes payment data each year on June 30 for the preceding calendar year, the initial implementation was different. Companies were first required to report payments beginning in August 2013, and CMS released this data on September 30, 2014, via the Open Payments Program website. As a result, data for 2014 may partially reflect both the pre- and post-transparency regimes.

This does not pose a threat to our empirical strategy. If anything, it implies that our DiD estimates understate the true magnitude of the policy's effects, without affecting their sign. Moreover, it is plausible that pharmaceutical firms did not adjust their strategies immediately following the first disclosure of data. Coordination on new equilibria may have taken time, suggesting a gradual transition occurred. If this is so, our estimates again represent a conservative lower bound of the policy's true impact—one that would have been fully captured by the DiD analysis had the 2014 data only reflected the pre-transparency regime and the post-transparency equilibrium been fully established by 2015.[39]

To implement our DiD analysis, we estimate a series of regression models that control for both observed and unobserved factors potentially affecting the outcome variables. The general form of the model is:

$$Y_{idt} = \delta Post_t \times Disclosure_d + \tau_t + \varphi_d + \varepsilon_{idt}, \tag{19}$$

where Y denotes the outcome variable of interest, such as the prescriptions made by physicians, the value of payments received, or some function thereof. Subscripts d, i, and t denote physicians, drugs, and time, respectively. The specification includes physician fixed effects φ_d and time fixed effects τ_t, which flexibly account for unobserved heterogeneity across doctors and time trends. In alternative specifications, we also include doctor–drug fixed intercepts $\varphi_d \times \gamma_i$ to absorb all time-invariant variation at the physician–drug level.

Our main parameter of interest is δ, which captures the effect of the introduction of transparency. We estimate this effect for the entire post-transparency period as well as separately for each individual year.

It is important to note that the identifying variation differs across specifications. In the baseline models with physician fixed effects φ_d, the relevant variation comes from differences in changes pre- and post-disclosure, as well as variations across drugs. In contrast, in specifications that use doctor-drug intercepts $\varphi_d \times \gamma_i$, the identifying variation comes solely

institutions that already apply to certain areas or groups are extended to new areas or groups. For instance, Kotchen and Grant (2011) examine the effect of daylight saving time on energy consumption by leveraging its extension to all counties in Indiana, while Rossi and Villar (2020) analyze the impact of extending subsidies for anti-malaria products to a broader set of households.

[39] One might argue that an alternative control group could consist of physicians who received no payments in 2014 or in subsequent years, and who therefore might be considered unaffected by the transparency regulations. However, selection into this group is clearly endogenous. Moreover, Agha and Zeltzer (2022) show that physicians can be influenced not only directly through payments from pharmaceutical companies, but also indirectly through interactions with colleagues who do receive such payments. While our theoretical model does not explicitly incorporate these peer effects, they may be relevant in practice. As a result, even physicians who receive no payments may alter their behavior in response to mandatory disclosure, rendering them unsuitable as a valid control group.

from differential pre–post changes over time for the same doctor–drug pair, eliminating cross-drug variation as a source of identification.

6.3 Results

We now employ the aforementioned approach to assess the impact of mandatory disclosure on prescription volumes and polarization, payment levels and polarization, and the marginal rate of return on these payments.

6.3.1 Level of prescriptions

We begin by examining prescription levels. Our model predicts that transparency reduces x, the share of patients treated with branded drugs, thereby lowering the total dollar value of prescriptions.

Table 3: Prescription volumes

Dep. variable: Prescriptons	**Full sample**		**Payments = 0**		**Payments > 0**	
	(1)	(2)	(3)	(4)	(5)	(6)
Transparency × Post	3449.744		5745.245		-13733.627**	
	(3096.392)		(3862.044)		(5981.756)	
Transparency × $Post_{2015}$		4038.961***		5374.327**		-2159.234
		(1529.114)		(2460.271)		(3330.166)
Transparency × $Post_{2016}$		2586.620		5439.730*		-9662.126*
		(2544.793)		(3184.485)		(5400.849)
Transparency × $Post_{2017}$		3750.638		6686.697		-17092.306**
		(3619.439)		(4376.758)		(7560.597)
Transparency × $Post_{2018}$		3424.036		5487.615		-32450.812***
		(5216.107)		(6225.331)		(11210.320)
Mean dep. var.	[110534.61]	[110534.61]	[92353.35]	[92353.35]	[121179.08]	[121179.08]
Year FE	YES	YES	YES	YES	YES	YES
Doctor FE	YES	YES	YES	YES	YES	YES
R^2	0.846	0.846	0.866	0.866	0.864	0.864
Observations	79882	79882	29498	29498	50384	50384

Notes: *Prescriptions* is the yearly total amount of NOACC prescribed by the doctor from all manufacturer; *Transparency × Post* is the interaction between the indicator for transparency regulation and the post-transparency period; *Transparency × $Post_{year}$* are the interactions between the indicator for transparency regulation and the single post-transparency years. Standard errors are clustered at the physician level. *, **, and *** denote significance at the 10, 5, and 1 percent level, respectively.

Table 3 reports the results of the DiD regression for the full sample and separately for physicians who receive payments. In the full sample, the effect of transparency is not statistically significant. Among paid physicians, however, the effect is both statistically significant and, in line with our theoretical predictions, negative. Although initially modest, the effect grows over time—consistent with a gradual adjustment process—and culminates in a sizeable reduction of more than 20% relative to the initial prescription volume in the final year of the study period.

6.3.2 Level of payments

We next examine the impact of transparency on the level of industry payments. In this case, the predictions of our model are ambiguous: while transparency alters the nature of competition for influence, this competition may either intensify or weaken. This ambiguity is illustrated in Figure 5, which shows the equilibrium level of payments for each physician type.

Table 19 presents the results of the DiD regression for both the full sample and the subsample of physicians who receive positive payments. Transparency appears to increase the level of payments, but the effect is not statistically significant and shows no clear trend toward significance over time. This result is consistent with the model's prediction of an ambiguous overall effect.

Table 4: Payment levels

Dep. variable: Payments	**Full sample**		**Payments > 0**	
	(1)	(2)	(3)	(4)
Transparency × Post	261.531		813.197	
	(303.720)		(1501.919)	
Transparency × Post_{2015}		228.062		627.290
		(216.996)		(978.368)
Transparency × Post_{2016}		449.001		1408.896
		(323.296)		(1593.876)
Transparency × Post_{2017}		168.083		269.534
		(441.876)		(2407.875)
Transparency × Post_{2018}		137.977		954.379
		(433.602)		(2841.315)
Mean dep. var.	[947.37]	[947.37]	[1502.03]	[1502.03]
Year FE	YES	YES	YES	YES
Doctor FE	YES	YES	YES	YES
R^2	0.754	0.754	0.761	0.761
Observations	79882	79882	50384	50384

Notes: *Payments* is the yearly total amount of payments received by the doctor from all manufacturer; *Transparency* × *Post* is the interaction between the indicator for transparency regulation and the post-transparency period; *Transparency* × $Post_{year}$ are the interactions between the indicator for transparency regulation and the single post-transparency years. In columns (1) and (2) we employ the full sample while in column (3) and (4) we restrict to doctors receiving payments. Standard errors are clustered at the physician level. *, **, and *** denote significance at the 10, 5, and 1 percent level, respectively.

6.3.3 Polarization

We now turn to the most distinctive implication of our theory: the polarization of both payments and prescriptions. Our model predicts that under mandatory disclosure, fewer doctors will prescribe both branded drugs and receive positive contributions from both pharmaceutical companies [40] Furthermore, even those doctors who continue to prescribe both drugs will be more heavily influenced by the nearest pharmaceutical company, resulting in more extreme choices.

Starting with prescriptions, we consider two alternative measures of polarization: the Herfindahl-Hirschman Index (HHI) of prescription volumes, and a count of the number of different drugs prescribed by a single doctor. The results are reported in Table 5.

Table 5: Polarization of prescriptions

Dep. variable:	**HHI prescriptions**				**Num. drugs prescribed**	
	Full sample		**Payments > 0**		**Full sample**	
	(1)	(2)	(3)	(4)	(5)	(6)
Transparency × Post	815.459***		471.161**		-0.285***	
	(109.153)		(186.700)		(0.040)	
Transparency × Post_{2015}		414.553***		454.893**		-0.135***
		(110.907)		(179.660)		(0.041)
Transparency × Post_{2016}		647.221***		379.807*		-0.262***
		(126.330)		(217.567)		(0.046)
Transparency × Post_{2017}		1001.385***		557.289***		-0.355***
		(124.962)		(213.641)		(0.047)
Transparency × Post_{2018}		1178.083***		515.638**		-0.382***
		(128.371)		(229.058)		(0.048)
Mean dep. var.	[5282.22]	[5282.22]	[5079.13]	[5079.13]	[3.31]	[3.31]
Year FE	YES	YES	YES	YES	YES	YES
Doctor FE	YES	YES	YES	YES	YES	YES
R^2	0.599	0.600	0.625	0.625	0.692	0.692
Observations	79882	79882	50384	50384	79882	79882

Notes: *HHI prescriptions* is the HHI index of concentration calculated for prescriptions; *Num. drugs prescribed* is the number of different NOACs prescribed at least once by the doctor in the current year; *Transparency* × *Post* is the interaction between the indicator for transparency regulation and the post-transparency period; *Transparency* × $Post_{year}$ are the interactions between the indicator for transparency regulation and the single post-transparency years. Columns (1) and (2) and columns (8) and (9) make use of the full sample while columns (3) to (7) restrict the sample to doctors receiving a yearly total of at least $50 to $1000. Standard errors are clustered at the physician level. *, **, and *** denote significance at the 10, 5, and 1 percent level, respectively.

The table provides strong evidence of a substantial polarization effect of mandatory disclosure on prescribing behavior. When we restrict the sample to doctors receiving positive payments, the HHI of prescriptions increases by nearly 500 points following the introduction of transparency, an increase of almost 10% from the initial level of approximately 5,100. As before, the effect appears to intensify over time, consistent with a gradual transition to a new equilibrium.[41]

[40] This follows from the fact that corner solutions where $y_1 = 0$ or $y_2 = 0$ become more prevalent, as illustrated in Figure 3).

[41] This time, the effect is even more pronounced in the full sample, where the HHI increases by over 15%

The evidence of polarization is further supported by the analysis of the number of distinct drugs prescribed, shown in the last column of Table 5. Under mandatory disclosure, fewer physicians prescribe multiple branded drugs. This effect is strongly significant and, once again, grows stronger over time.

We adopt a similar approach in analyzing the polarization of payments. Specifically, we examine both the HHI of payments and the number of pharmaceutical companies from which each doctor receives payments. Naturally, this analysis is limited to physicians who receive positive payments. However, we also explore more restricted subsamples defined by payment thresholds—specifically, doctors receiving at least \$500 or \$1,000 annually (note that \$500 is less than one-third of the average annual payment in our sample).

In the full sample (first two columns of Table 6), the HHI of payments increases following the introduction of transparency, and this effect appears to strengthen over time. However, the estimate is not statistically significant, possibly because the full sample includes physicians who receive very small payments, rendering the HHI a less reliable measure of concentration.[42] When we restrict the analysis to physicians receiving at least \$500 or \$1,000 per year, the effect of transparency on payment concentration becomes stronger, reaching 15% of the initial level, and statistically significant. However, the increase in the HHI now becomes more erratic across time periods, likely due to the reduced sample size in these higher-payment groups.

The estimates of the number of manufacturers making payments to a single physician, reported in the last two columns of Table 6, indicate that transparency leads to a reduction in the number of paying firms. The effect is statistically significant and, once again, becomes more pronounced over time.

6.3.4 Marginal product of contributions

Running the regression described in Subsection 6.2 separately for the pre- and post-transparency periods reveals a substantial increase in the estimated coefficient, which nearly doubles from approximately 1 to almost 2 (see Table 7). It is tempting to interpret this coefficient as the marginal product of contributions; that is, the increase in the value of a physician's branded drug prescriptions associated with a one-dollar increase in payments from the drug manufacturer.

At first glance, the sharp and statistically significant increase of such marginal product may seem surprising. Upon closer consideration, however, it is fully consistent with our theoretical model. The intuition is straightforward. In the common-influence equilibrium that prevails in the absence of transparency, contributions from competing firms tend to cancel each other out in terms of their influence on a physician's prescribing behavior. By

relative to the pre-disclosure level. While our model, taken literally, predicts no effect of transparency on physicians who do not receive payments, peer effects may nonetheless cause even these doctors to adjust their prescribing behavior in response to the policy, as discussed in the previous footnote.

[42] Low payment levels may generate irregular payment patterns, whereby doctors receive small, sporadic payments from different companies across years. This can result in substantial year-to-year variation in the HHI that does not meaningfully reflect shifts in the underlying concentration of payments.

Table 6: Polarization of payments

Dep. variable:	HHI payments						Num. payers	
	Payments > 0 every year		**Payments > 500**		**Payments > 1000**		**Full sample**	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Transparency × Post	269.596		969.697***		1213.589***		-0.140***	
	(305.781)		(339.659)		(394.435)		(0.034)	
Transparency × $Post_{2015}$		161.610		1198.433***		1376.950***		-0.046
		(421.061)		(400.499)		(466.429)		(0.036)
Transparency × $Post_{2016}$		150.726		591.612		794.795		-0.142***
		(378.743)		(429.242)		(549.948)		(0.040)
Transparency × $Post_{2017}$		410.850		1266.158		1636.739**		-0.225***
		(418.238)		(845.630)		(777.152)		(0.041)
Transparency × $Post_{2018}$		746.894		654.038		831.913*		-0.315***
		(473.604)		(519.871)		(454.226)		(0.038)
Mean dep. var.	[6064.04]	[6064.04]	[6636.42]	[6636.42]	[8438.73]	[8438.73]	[1.39]	[1.39]
Year FE	YES	YES	YES	YES	YES	YES	YES	YES
Doctor FE	YES	YES	YES	YES	YES	YES	YES	YES
R^2	0.567	0.567	0.898	0.898	0.842	0.843	0.809	0.810
Observations	36612	36612	7250	7250	3744	3744	79882	79882

Notes: *HHI payments* is the HHI index of concentration calculated for payments; *Num. payers* is the number of different manufacturers (from 1 to 3) that paid a positive amount for detailing to the doctor in the current year; *Transparency × Post* is the interaction between the indicator for transparency regulation and the post-transparency period; *Transparency* × $Post_{year}$ are the interactions between the indicator for transparency regulation and the single post-transparency years. Columns (1) and (2) and columns (8) and (9) make use of the full sample while columns (3) and (4) restrict the sample to doctors receiving a yearly total of at least $500 while columns (5) and (6) to doctors receiving a yearly total of at least $1000. Standard errors are clustered at the physician level. *, **, and *** denote significance at the 10, 5, and 1 percent level, respectively.

contrast, in the prevalent-influence equilibria arising under transparency, a single firm—the primary contributor—exerts influence on a doctor's market share choices. The absence of countervailing influence from rivals amplifies the effect of each dollar contributed, resulting in a greater impact on prescribing behavior. This pattern thus represents a distinctive and testable implication of our model, corroborated by the data.

Table 7: Marginal returns on payments

Dependent variable: Amount prescribed		**All years**		**Pre-trasparency**	**Post-trasparency**
	(1)	(2)	(3)	(4)	(5)
Amount received	2.452***	2.268***	1.721***	1.021***	1.954***
	(0.183)	(0.170)	(0.144)	(0.202)	(0.143)
Year FE	YES	YES	YES	YES	YES
Doc FE	NO	YES	YES	YES	YES
Drug FE	NO	NO	YES	YES	YES
R^2	0.080	0.302	0.488	0.553	0.528
Observations	319528	319528	319528	59660	259868

Notes: *Amount prescribed* is the yearly total amount of NOACC prescribed by the doctor from a single manufacturer; *Amount received* is the yearly total amount of payments received by the doctor from the same manufacturer. Columns (1) to (3) make use of the full sample while columns (4) and (5) focus on the pre-transparency period and the post-transparency period respectively. Standard errors are clustered at the physician level. *, **, and *** denote significance at the 10, 5, and 1 percent level, respectively.

7 Structural estimation

This section describes the equilibrium model of payments and prescriptions that we estimate in order to compute the welfare changes associated to the introduction of transparency in the market for NOACs.

7.1 Empirical model of duopoly of branded drugs and the generic

The model developed in Sections 2, 3, and 4, derives predictions on equilibrium outcomes of the competition between two principals selling their own product (two manufacturers of branded drugs) when the agent also has the possibility to allocate part of the demand to a third product (a generic drug). Although an extension to three principals is discussed in Section E in the Appendix, we focus here on the two main branded drug manufactures, Eliquis and Xarelto, and we consider the prescriptions of other drugs other than these two as the alternative choice of the agent. Thus, the choice x is going to be the combined share of these two branded drugs in doctors' prescriptions, while the choice y is going to be split of such combined market share between the two branded drugs.

The empirical model that we estimate relies on the information on contributions made by the branded firms (the manufacturers of the generic do not report any contribution),[43] and on the prescriptions made by doctors. Our theoretical model provides equilibrium relations between these variables, depending on the transparency regime, which we estimate using data before transparency is introduced in 2014, and after transparency has been put in

[43]This is consistent with the fact that many pharmaceutical companies produce the generic, margins are relatively low, and the generic is considered as homogeneous, thus eliminating the incentives for these manufacturers to invest resources in detailing.

place in 2018.[44]

As shown in our theoretical model, the equilibrium choices of y and x made by doctors, and the equilibrium transfers T made by firms, are determined by the transparency regime $TR = \{NT, T\}$, by the preference parameters θ and ξ, by the dispersion Δ of the distribution of doctors' preference parameter θ, by the resistance parameters μ and ν, by the profit margin π, and by the equilibrium that is selected, when multiple equilibria are possible. In particular, there are two possible equilibria under the regine of no-transparency: weak influence and strong influcence – denoted with W and S respectively; and a continuum of equilibria in the regime of transparency, which are intermediate between the two extreme cases of exclusivity and the most efficient equilibrium of prevalent influence – denoted with E and P respectively –. We make some assumptions on the distributions of the parameters and on what firms can observe. First, we assume that, consistently with our theoretical model, θ is uniformly distributed around an average value of m. We thus model $\theta \sim U[\theta_m, \theta_\Delta]$ where θ_m is the average of the uniform distribution and θ_Δ is the length of its support. Second, we assume that the remaining parameters related to the doctors' preferences and resistance are normally distributed – $u \sim N(\mu_u, \sigma_u)$ where u denotes the parameter of interest and μ and σ are the mean of standard deviation of the distribution –, and that firms can observe such characteristics. Third, we assume that the profit margin for the branded products is common knowledge and the same for the two firms. Fourth, we assume that under transparency only the two equlibria E and P are possible. Thus, the goal of the estimation is to recover the following set of parameters: $\Theta = (\theta_m, \theta_\Delta, \mu^\mu, \sigma^\mu, \mu^\nu, \sigma^\nu, \mu^\xi, \sigma^\xi, \pi)$.

A. Prescriptions

The observed share of prescriptions of branded products y made by a doctor i in our sample in the two transparency regimes can be written compactly as in equations 20.

$$\begin{cases} y_i^{NT,W} = & y_i^{NT,S} = 3\theta_m - 1 \\ y_i^{T,E} = & y_i^{T,P} = y_i^M(\Theta) \end{cases} \tag{20}$$

As described in Sections 3 and 4.2, the equilibrium prescriptions without transparency are the same in the two equilibria of weak and strong influence; similarly, with transparency and exclusive influence $y^{T,E}$ are the same as in the case of prevalent influence $y^{T,P}$.

The observed

8 Conclusions

This paper presents a novel theory that sheds light on how transparency regulations impact competition for influence and provides empirical evidence of its implications in the market

[44]We select this year for the post-transparency regime so to give the longest amount of time in our data for the new equilibrium to arise. This choice is also supported by the reduced form analysis in Section 6 which shows a progressive increase in polarization throughout the post-treatment years.

for Novel Oral Anti-Coagulants (NOACs). Our theory explains how mandatory disclosure of payments made to decision makers (such as doctors, politicians, and investment managers) alter the dynamics of competition among principals (such as special interest groups, pharmaceutical companies, and issuers of financial assets). While transparency regulations aim to reduce potential conflicts of interest, we argue that they also create distinct spheres of influence, in which decision makers are primarily influenced by a single principal. This, in turn, eliminates the opposing influence of competing principals, leading to polarization of payments and choices.

We document the polarization effect in the market for NOACs by analyzing data from a period before and after the implementation of the Sunshine Act of 2010. We compare states with and without transparency regulations before the act and employ a difference-in-differences analysis to identify the effects.

The empirical analysis provides strong support for our theoretical predictions. Specifically, we find robust evidence of polarization, consistent with the formation of distinct spheres of influence among decision makers following the introduction of mandatory disclosure. Furthermore, we confirm that transparency significantly enhances the estimated productivity of payments, due to the elimination of opposing influence from rivals within each pharmaceutical company's respective sphere of influence.

Our theory implies that mandatory disclosure can have both positive and negative effects on social welfare. On the one hand, the polarization of prescriptions resulting from transparency may lead to more patients being treated with branded drugs that are not the optimal choice for their individual health needs. On the other hand, transparency may also lead to fewer patients being prescribed branded drugs when an unbranded option would be more appropriate. To determine which effect dominates in practice, future research should undertake a structural estimation of our model. Such an analysis would provide a comprehensive assessment of the impact of transparency on social welfare that takes into account all the factors that are relevant in our model.

[Tra le estensioni, discutere il caso in cui i medici sono connessi in un network e i lavori su influencing connected politicians di Battaglini e Petacchini JPE 2018]

References

Agha, L. and Zeltzer, D. (2022). Drug diffusion through peer networks: The influence of industry payments. *American Economic Journal: Economic Policy*, 14(2):1–33.

Allcott, H. and Taubinsky, D. (2015). Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market. *American Economic Review*, 105(8):2501–2538.

Altonji, J. G., Conley, T., Elder, T. E., and Taber, C. R. (2010). Methods for using selection on observed variables to address selection on unobserved variables. *Mimeo.*

Altonji, J. G., Elder, T. E., and Taber, C. R. (2005). Selection on observed and unobserved variables: Assessing the effectiveness of catholic schools. *Journal of political economy*, 113(1):151–184.

Amaral-Garcia, S. (2022). Medical device companies and doctors: Do their interactions affect medical treatments? *University of York, Health, Econometrics and Data Group (HEDG) Working Paper 22/10.*

Barber, M., McCarty, N., Mansbridge, J., and Martin, C. J. (2015). Causes and consequences of polarization. *Political negotiation: A handbook*, 37:39–43.

Bernheim, B. D. and Whinston, M. D. (1986). Menu auctions, resource allocation, and economic influence. *The quarterly journal of economics*, 101(1):1–31.

Bernheim, B. D. and Whinston, M. D. (1998). Exclusive dealing. *Journal of political Economy*, 106(1):64–103.

Blumenthal, D. (2004). Doctors and Drug Companies. *New England Journal of Medicine*, 351(18):1885–1890.

Bronnenberg, B. J., Dubé, J.-P., Gentzkow, M., and Shapiro, J. M. (2015). Do pharmacists buy bayer? informed shoppers and the brand premium. *The Quarterly Journal of Economics*, 130(4):1669–1726.

Cain, D. M., Loewenstein, G., and Moore, D. A. (2005). The dirt on coming clean: Perverse effects of disclosing conflicts of interest. *The Journal of Legal Studies*, 34(1):1–25.

Carey, C., Lieber, E., and Miller, S. (2021). Drug Firms' Payments and Physicians' Prescribing Behavior in Medicare Part D. *Journal of Public Economics*, 26751(197):104402.

Chao, M. and Larkin, I. (2022). Regulating conflicts of interest in medicine through public disclosure: Evidence from a physician payments sunshine law. *Management Science*, 68(2):1078–1094.

Chen, D. L., Levonyan, V., Reinhart, S. E., and Taksler, G. (2019). Mandatory disclosure: Theory and evidence from industry-physician relationships. *The Journal of Legal Studies*, 48(2):409–440.

Christensen, H. B., Floyd, E., and Maffett, M. (2020). The only prescription is transparency: The effect of charge-price-transparency regulation on healthcare prices. *Management Science*, 66(7):2861–2882.

Ershov, D. and Mitchell, M. (2022). The effects of influencer advertising disclosure regulations: Evidence from instagram. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 73–74.

Fainmesser, I. P. and Galeotti, A. (2021). The market for online influence. *American Economic Journal: Microeconomics*, 13(4):332–72.

Fiorina, M. P., Abrams, S. A., and Pope, J. C. (2008). Polarization in the american public: Misconceptions and misreadings. *The Journal of Politics*, 70(2):556–560.

Grennan, J. and Michaely, R. (2021). Fintechs and the market for financial analysis. *Journal of Financial and Quantitative Analysis*, 56(6):1877–1907.

Grennan, M., Myers, K., Swanson, A., and Chatterji, A. (2021). No free lunch? welfare analysis of firms selling through expert intermediaries. *NBER Working Paper*.

Grossman, G. M. and Helpman, E. (1994). Protection for sale. *The American Economic Review*, 84(4):833.

Grossman, G. M. and Helpman, E. (2001). *Special interest politics*. MIT press.

Guo, T., Sriram, S., and Manchanda, P. (2021). The effect of information disclosure on industry payments to physicians. *Journal of Marketing Research*, 1(58):115–140.

Handel, B. R. and Kolstad, J. T. (2015). Health insurance for "humans": Information frictions, plan choice, and consumer welfare. *American Economic Review*, 105(8):2449–2500.

Inderst, R. and Ottaviani, M. (2009). Misselling through agents. *American Economic Review*, 99(3):883–908.

Iyer, S., Yoo, J. S., Jenkins, N. W., Parrish, J. M., Khechen, B., Haws, B. E., Nolte, M. T., Brundage, T. S., Hrynewycz, N. M., Phillips, F. M., et al. (2020). All disclosure is good disclosure: patient awareness of the sunshine act and perceptions of surgeon-industry relationships. *Clinical Spine Surgery*, 33(3):96–100.

Jensen, J., Naidu, S., Kaplan, E., Wilse-Samson, L., Gergen, D., Zuckerman, M., and Spirling, A. (2012). Political polarization and the dynamics of political language: Evidence from 130 years of partisan speech [with comments and discussion]. *Brookings Papers on Economic Activity*, pages 1–81.

Kjerpeseth, L. J., Selmer, R., Ariansen, I., Øystein Karlstad, Ellekjær, H., and Skovlund, E. (2019). Comparative effectiveness of warfarin, dabigatran, rivaroxaban and apixaban in non-valvular atrial fibrillation: A nationwide pharmacoepidemiological study. *PLoS One*, 14(8):e0221500.

Kotchen, M. J. and Grant, L. E. (2011). Does daylight saving time save energy? evidence from a natural experiment in indiana. *Review of Economics and Statistics*, 93(4):1172–1185.

Martimort, D. and Semenov, A. (2008a). Ideological uncertainty and lobbying competition. *Journal of Public Economics*, 92(3-4):456–481.

Martimort, D. and Semenov, A. (2008b). The informational effects of competition and collusion in legislative politics. *Journal of Public Economics*, 92(7):1541–1563.

Martimort, D. and Stole, L. (2015). Menu auctions and influence games with private information. *Available at SSRN 2569703*.

McCarty, N., Poole, K. T., and Rosenthal, H. (2016). *Polarized America: The dance of ideology and unequal riches*. MIT Press.

Mitchell, M. (2021). Free ad(vice): internet influencers and disclosure regulation. *The RAND Journal of Economics*, 52(1):3–21.

O'Brien, D. P. and Shaffer, G. (1997). Nonlinear supply contracts, exclusive dealing, and equilibrium market foreclosure. *Journal of Economics & Management Strategy*, 6(4):755–785.

Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, 37(2):187–204.

Pham-Kanter, G., Mello, M. M., Lehmann, L. S., Campbell, E. G., and Carpenter, D. (2017). Public awareness of and contact with physicians who receive industry payments: a national survey. *Journal of general internal medicine*, 32(7):767–774.

Poole, K. T. and Rosenthal, H. (1984). The polarization of american politics. *The journal of politics*, 46(4):1061–1079.

Robertson, C., Rose, S., and Kesselheim, A. S. (2012). Effect of financial relationships on the behaviors of health care professionals: a review of the evidence. *The Journal of Law, Medicine and Ethics*, 40(3):452–466.

Rossi, P. and Villar, P. (2020). Private health investments under competing risks: Evidence from malaria control in senegal. *Journal of Health Economics*, 73:102330.

Safavi, K. S., Hong, A., Janney, C. F., Panchbhavi, V. K., and Jupiter, D. C. (2022). Sunshine act in the dark. *Clinical Ethics*, 17(2):122–129.

Shah, S. V. and Gage, B. F. (2011). Cost-Effectiveness of Dabigatran for Stroke Prophylaxis in Atrial Fibrillation. *Circulation*, 123(22):2562–2570.

Sinha, M. S. and Kesselheim, A. S. (2017). The effects of the sunshine act: What can and should we expect? *The American Journal of Bioethics*, 17(6):22–24. PMID: 28537826.

Thacker, P. D., Kesselheim, A. S., and Campbell, E. G. (2014). Will a new website empower patients to ask their physicians about financial relationships with industry? In *JAMA Forum Archive*, volume 3. American Medical Association.

Yao, X., Abraham, N. S., Sangaralingham, L. R., Bellolio, M. F., McBane, R. D., Shah, N. D., and Noseworthy, P. A. (2016). Effectiveness and Safety of Dabigatran, Rivaroxaban, and Apixaban Versus Warfarin in Nonvalvular Atrial Fibrillation. *Journal of the American Heart Association*, 5(6):e003725.

A Complete information

Under complete information and without mandatory disclosure, the model becomes a special case of Bernheim and Whinston (1986) and Grossman and Helpman (1994). As is well known, this common agency game has multiple equilibria. To address this multiplicity, the literature has generally focused on *truthful* equilibria, where each principal offers a payment that reflects its true payoff. Such equilibria are selected for various reasons, as discussed in Bernheim and Whinston (1986) and Grossman and Helpman (1994). In particular, truthful equilibria are Pareto dominant for the two principals and lead to efficient actions, meaning actions that maximize the players' joint surplus.

To analytically present the truthful equilibrium of our model, it is useful to introduce some further notation. For any coalition of the agent and a subset $S \subseteq \{1, 2\}$ of principals, define the joint payoff of the coalition as:

$$V_S = U_0 + \sum_{i \in S} U_i(x, y),$$

and the jointly optimal action as:

$$(x^*_S, y^*_S) = \arg \max_{0 \leq x, y \leq 1} V_S.$$

The associated joint payoff is denoted by $V^*_S = V_S(x^*_S, y^*_S)$, and it can be thought of as the characteristic function of the associated cooperative game.

With this notation, the agent's action under complete information is:

$$\left(x^*_{\{1,2\}}, y^*_{\{1,2\}} \right) = \left(\min \left[\xi + \frac{\pi}{2\nu}, 1 \right], \theta \right).$$

Thus, the complete information equilibrium exhibits over-treatment with respect to the doctor's bliss point, $x^*_{\{1,2\}} > \xi$, to a degree that depends on the ratio between π and ν. Instead, the choice of y coincides with the doctor's bliss point for y, because the principals influence y in opposite directions, and they offset each other's efforts exactly. Hence, although each principal has a unilateral incentive to pay the agent to steer his choice along the y-dimension towards its preferred direction, these payments are ultimately wasteful.

The truthful contribution schedules are:

$$\Gamma_i(x, y_i) = -F_i + \pi x y_i,$$

and thus are linear in x given y, and linear in y given x. Principals' net payoffs coincide with their respective F_i. If the characteristic function is sub-additive, i.e. $V^*_{\{1,2\}} \leq V^*_{\{1\}} + V^*_{\{2\}}$, these constant terms are $F_i = \left(V^*_{\{1,2\}} - V^*_{\{j\}} \right)$ and thus coincide with each principal's marginal contributions to the grand coalition. When instead the characteristic function is super-additive, i.e., $V^*_{\{1,2\}} > V^*_{\{1\}} + V^*_{\{2\}}$, there is a multiplicity of equilibria where $F_1 + F_2 = V^*_{\{1,2\}}$ so that the agent's payoff vanishes, but the division of the social surplus between the principals is indeterminate.

Under complete information, transparency can give rise to different equilibria.[45] However, the set of Pareto-undominated equilibria exhibits identical actions and payoffs with and without transparency.

This is obvious when the truthful equilibrium entails zero rent for the agent. When instead in the truthful equilibrium the agent obtains a positive rent, and principals earn their marginal contribution, the result follows from a general property of common agency games with complete information, known as bilateral efficiency. Bilateral efficiency means that, with and without transparency, each principal's contribution schedule must maximize the joint payoff of the agent and that principal, given the other principal's schedule. (For a formal proof, see Bernheim and Whinston (1986).) This implies that neither principal can receive a payoff greater than its marginal contribution, or it would be excluded by the coalition of the agent and the other principal. Since the truthful equilibrium yields each principal's marginal contribution as its payoff, it cannot be improved upon even if the space of feasible strategies is larger. Therefore, in Pareto-undominated equilibria, equilibrium actions and payoffs must be the same with and without transparency.

B Non-linearity of contribution schedules in the NOAC market

[TO BE ADDED]

C Equilibrium schedules

This appendix provides the equilibrium contribution schedules omitted from the main text.

In all scenarios—monopolistic influence, common influence, and prevalent influence—the equilibrium schedules are piecewise quadratic in x and y. Each quadratic segment takes the form:

$$\alpha_{0i} + \alpha_{1i} y_i + \alpha_{2i} y_i^2 + \alpha_{3i} x + \alpha_{4i} x^2 + \alpha_{5i} x y_i = \alpha_{0i} + \Gamma_i(x, y_i), \tag{21}$$

where $\Gamma_i(x, y_i)$ denotes the variable part of the schedule.

Since we focus on the symmetric case where $\pi_1 = \pi_2 = \pi$ and $\Lambda = \frac{1}{2}$, we report only principal 1's schedules. The equilibrium schedules for principal 2 mirror those of principal 1, with y replaced by $1 - y$.

[45]For example, principals may offer positive contributions only on the condition that the agent does not receive any contribution from the rival. Such strategies are self-sustaining: if both principals use contribution schedules of this form, neither has a unilateral incentive to deviate. In the resulting exclusive-influence equilibrium, the agent's choice is either $\left(x^*_{\{1\}}, y^*_{\{1\}}\right)$ or $\left(x^*_{\{2\}}, y^*_{\{2\}}\right)$, and principal i's payoff is $\pi x^*_{\{j\}} y^*_{\{j\}}$. In other words, principals compete in a Bertrand-like fashion for the exclusive right to influence the agent's choice, and in equilibrium, they are indifferent between winning or losing the competition.

C.1 Monopolistic influence

Under monopolistic influence by principal 1, assuming that $4\mu\nu > \pi^2$ to avoid a degenerate solution, the equilibrium payment schedule is:

$$\Gamma_1(x,y) = \begin{cases} -\frac{(1+\Delta)\mu+\xi\pi}{2}y + \frac{4\mu\nu-\pi^2}{8\nu}y^2 + \pi xy \quad (\equiv \Gamma_1^M(x,y)) & \text{if } y \geq \frac{2\nu[(1+\Delta)\mu-\xi\pi]}{\pi^2+4\mu\nu}, \\ -\xi\pi y - \frac{\pi^2}{4\nu}y^2 + \pi xy & \text{if } y \leq \frac{2\nu[(1+\Delta)\mu-\xi\pi]}{\pi^2+4\mu\nu}. \end{cases}$$

The schedule comprises two branches, depending on whether $y^{M_1}(\theta)$ is greater than or equal to θ – a condition that is equivalent to y being greater or lower than $\frac{2\mu[(1+\Delta)\mu-\xi\pi]}{\pi^2+4\mu^2}$. In the upper branch, where we have proper monopolistic influence, the principal influences both x and y, while in the lower branch, the principal influences only his choice of x. This can be seen by noting that the derivative $\frac{\partial\Gamma_1(x,y)}{\partial y}$ vanishes at the equilibrium action $y = \theta$.

When $\Delta < \frac{\pi(4\xi\nu+\pi)}{8\mu\nu+\pi^2}$, only the upper branch of the schedule is relevant, and the constant term α_0 guarantees that the type $\theta = \theta_{\min}$ is indifferent between accepting principal 1's or not. When instead $\Delta > \frac{\pi(4\xi\nu+\pi)}{8\mu\nu+\pi^2}$, whereas α_0 guarantees that the schedule is continuous at the switching point between the lower and upper branches. The explicit expressions for α_0 can be found in the online mathematical appendix.

C.2 Common influence

As discussed in the main text, ruling out cases where the incentive to influence the agent's choice of y is so weak that principals do not seek to influence the middle type, two forms of equilibrium can arise, corresponding to the case of strong and weak competition for influence. In the case of strong competition, both principals influence the choice of y of all types, whereas in the case of weak competition each principals influences the choice of y of its closest types only.

Thresholds for strong and weak competition. We start by stating the conditions under which each of these cases arises.

The case of weak competition can only arise when the slope of the curve $y^{M_1}(\theta)$ is lower than the slope of the curve $y = 3\theta - 1$, which is 3, and that the two curves intersect at a value of θ greater than $\theta_{\min}$. These conditions hold when

$$\pi \leq \min\left[\frac{2\sqrt{\mu}\sqrt{\nu}}{\sqrt{3}}, \frac{2\left(\sqrt{\nu\left(6\Delta^2\mu + 2\Delta\mu + \xi^2\nu\right)} - \xi\nu\right)}{3\Delta + 1}\right]$$

When this inequality is reversed, the equilibrium is one of strong competition for influence.

Strong competition. Under strong competition, the equilibrium payment schedule is:

$$\Gamma_1^C(x,y) = -\frac{(1+3\Delta)\,\mu}{3}y + \frac{\mu}{3}y^2 + \pi xy \tag{22}$$

The constant term α_0 vanishes when $\Delta \leq 1/3$, implying that more extremist types choose either $y = 1$ or $y = 0$. When instead $\Delta > 1/3$, the constant α_0 is not uniquely determined. The reason for this is that in this case, the constant depends on those parts of the contribution schedules that apply for values of y that are not chosen in equilibrium and therefore are, to some extent, arbitrary. Following Martimort and Stole (2009), we focus on the *natural* equilibrium, which is obtained when the schedule reported above applies to all values of y, including those that are never chosen in equilibrium. The explicit expression is reported in the mathematical Appendix.

Weak competition. In this case, the equilibrium payment schedule comprises three branches:

$$\Gamma_1(x,y) = \begin{cases} -\frac{\pi(\pi+2\xi\nu)}{2\nu}y - \frac{\pi^2}{4\nu}y^2 + \pi xy & \text{for } y \leq 1-\bar{y} \\ \Gamma_1^C(x,y) & \text{for } 1-\bar{y} \leq y \leq \bar{y}, \\ \Gamma_1^M(x,y) & \text{for } y \geq \bar{y} \end{cases} \tag{23}$$

where $\bar{y} = \frac{2\upsilon(1-3\Delta)\mu+3\pi\xi}{4\mu\upsilon-3\pi^2}$. For low values of y, principal 1 elects not to influence the agent's choice of y. This branch of the schedule is similar to that arising under monopolistic influence, but now the value of $y(\theta)$ that principal 1 does not seek to change is $y^{M_2}(\theta)$ rather than θ. For intermediate values of y, the schedule takes the same form as under strong competition, as both principals influence the agent's choice of y. Finally, for high values of y, the schedule takes the same form as under proper monopolistic influence. This branch applies when only principal 1 influences y, which will therefore be equal to $y^{M_1}(\theta)$.

The constant terms of each branch are determined as follows. For the lower branch, α_0 guarantees that the type $\theta = \theta_{\min}$ is just indifferent between accepting principal 1's contribution or not, whereas the constants of the two upper branches guarantee that the schedule is continuous at the switching points between the different branches. The explicit expressions can be found in the mathematical appendix.

C.3 Transparency

Payment schedules with exclusive influence. The exclusive-influence contribution schedules are $\Gamma_1^E(x,y) = \Gamma_1^{M_1}(x,y)$. Thus, the only difference with the case of monopolistic influence is the constant term. The value of α_0 is reported in the mathematical appendix; it corresponds to the constant that applies under prevalent influence for $\bar{T} = 0$.

Payment schedules with prevalent influence. The contribution schedule comprises two branches. The upper branch, which applies when the principal is the primary contributor, is

$$\Gamma_1^H(x, y_i) = \Gamma_1^{M_1}(x,y), \tag{24}$$

whereas the lower branch, which applies to types in the sphere of prevalent influence of principal 2, is:

$$\Gamma_1^L(x,y) = \kappa\left[\frac{\pi^2}{2\upsilon}y^2 + 2\upsilon x^2 + 2\pi xy\right], \tag{25}$$

where

$$\kappa = \frac{\sqrt{\bar{T}\upsilon}}{(2\xi\upsilon + \pi) + 2\sqrt{\bar{T}\upsilon}}.$$

The constant terms of the schedules are reported in the mathematical appendix.

D Proofs

This appendix contains proof of the formal results presented in the main text. Some parts of certain proofs are only outlined; interested readers can find the complete details in the mathematical Appendix.

D.1 Proof of Proposition 1 (Monopolistic Influence)

To fix ideas, we focus on the case where the only principal that influences the agent is principal 1. The case of monopolistic influence by principal 2 can be solved similarly.

A monopolistic principal faces a screening problem that can be solved by invoking the Revelation Principle and thus focusing on direct mechanisms. Note that the payoff function $U_0(x, y, \theta)$ satisfies the sorting condition, as $\frac{\partial^2 U_0(x,y,\theta)}{\partial\theta\partial y} = 2\mu > 0$, and $\frac{\partial^2 U_0(x,y,\theta)}{\partial\theta\partial x} = 0$. Additionally, the agent's reservation payoff, which is achieved by setting $x = \xi$ and $y = \theta$, is zero. Standard arguments then imply that the solution can be found by pointwise maximization with respect to x and y of the virtual surplus function:

$$V = U_0(x, y, \theta) + U_1(x, y) - \frac{1 - F(\theta)}{f(\theta)} \frac{\partial U_0(x, y, \theta)}{\partial\theta},$$

provided that the resulting functions $\{x(\theta), y(\theta)\}$ are non-decreasing in θ – a condition that ensures incentive compatibility and will be checked later.

It is straightforward to verify that the virtual surplus is given by:

$$V(x, y, \theta) = -\nu(x - \xi)^2 - \mu(y - \theta)^2 + \pi xy - \mu (1 + \Delta - 2\theta) (y - \theta)$$

and thus is concave in (x, y). Therefore, we can focus on the first-order conditions for a maximum. Assuming full participation for the time being, $V(x, y, \theta)$ must be maximized for all values of θ. The solution for x is evidently:

$$x = \min \left[\xi + \frac{\pi}{2\nu} y, 1\right].$$

Regarding y, the first derivative is:

$$\frac{\partial V(x, y, \theta)}{\partial y} = -2\mu(y - \theta) + \pi x - \mu (1 + \Delta - 2\theta).$$

Clearly, if $\pi^2 \geq 4\mu\nu$ the derivative is always positive at $y = 1$, implying a corner solution $y(\theta) = 1$. If instead $\pi^2 < 4\mu\nu$, one obtains:

$$y = \min \left[2\nu \frac{\pi\xi - \mu(1 + \Delta)}{4\mu\nu - \pi^2} + \frac{8\mu\nu}{4\mu\nu - \pi^2}\theta, 1 \right]. \tag{26}$$

This is non-decreasing in θ, while $x(\theta)$ is constant, meaning that the solution is incentive compatible.

However, the assumption of full participation need not always hold. Specifically, while it is always optimal to incentivize all agent types in the x-dimension – where there is no asymmetry of information and, hence, no information rents – it may not be optimal to have full participation in the y-dimension. The intuitive reason is that including the more distant types in the scheme is costly, as it increases the information rents for the closer types.

To address this issue, we adapt techniques developed for uni-dimensional screening problems. In this case, full participation is optimal when the optimized virtual surplus is greater than what it would be if the agent maximized his own payoff, whereas partial participation is optimal when the optimized virtual surplus is lower.

In our problem, we can use the optimality condition $x = \max \left[\xi + \frac{\pi}{2\nu} y, 1 \right]$ to eliminate the variable x and transform the bi-dimensional screening problem into a uni-dimensional one. When the agent is not incentivized in the y-dimension, so that $y(\theta) = \theta$, but is optimally incentivized in the x-dimension, the virtual surplus becomes:

$$V(\theta) = \pi\xi\theta + \frac{\pi^2}{4\mu}\theta^2.$$

Therefore, full participation requires the optimized virtual surplus to be higher than this expression for all types θ. But if the above solution (26) yields a virtual surplus lower than this expression for some agent types, then for these types, the solution must entail $y(\theta) = \theta$. It is easy to verify that this occurs precisely when the solution (26) delivers a value of $y(\theta)$ lower than θ, that is, when $\theta < \frac{2\mu(\Delta\mu + \mu - \pi\xi)}{4\mu\nu - \pi^2}$. Therefore, the optimal value of $y(\theta)$ is given by condition (9). ■

This completes the proof of Proposition 1. Verifying that the solution we have characterized can be implemented using the contribution schedules presented in Appendix C is a straightforward exercise left to the reader.

D.2 Proof of Proposition 2 (Common Influence)

We begin by outlining the logic of the proof. As a first step, we show that, given the contribution schedules specified in Appendix C, the agent's optimal choices correspond to the equilibrium actions $\{x^C(\theta), y^C(\theta)\}$. We then examine the monopolistic screening problem faced by principal 1, taking as given the equilibrium schedule of principal 2. In this problem, the agent's indirect utility is the sum of his original payoff and the contribution received from principal 2. Applying the Revelation Principle, we solve the problem by considering direct mechanisms, in which principal 1 selects the actions $x(\theta)$ and $y(\theta)$ subject to the appropriate participation and incentive compatibility constraints. We demonstrate that the optimal actions chosen in this way coincide with the proposed equilibrium actions. Since these actions are implemented by the equilibrium schedules, it follows that principal 1's schedule constitutes a best response to that of principal 2. By symmetry, the same holds for principal 2, confirming that the specified schedules form an equilibrium of the influence game. Note that the use of direct mechanisms implies that the best response property holds in the full space of feasible schedules; in other words, we are not restricting principals to use quadratic payment schedules.

To fix ideas, we focus on the case of strong competition, where the variable components of the contribution schedules are:

$$\Gamma_1^C(x,y) = -\frac{(1+3\Delta)\,\mu}{3}y + \frac{\mu}{3}y^2 + \pi xy, \tag{27}$$

and

$$\Gamma_2^C(x,y) = -\frac{(1+3\Delta)\,\mu}{3}(1-y) + \frac{\mu}{3}(1-y)^2 + \pi x(1-y). \tag{28}$$

When confronted with these schedules, the agent chooses x and y to maximize $\Pi_0(x,y,\theta) = U_0(x,y,\theta) + \Gamma_1^C(x,y) + \Gamma_2^C(x,y)$. It is straightforward to verify that the function $\Pi_0(x,y,\theta)$ is concave, with first derivatives given by $\frac{\partial \Pi_0}{\partial x} = \pi - 2\nu\,(x-\xi)$ and $\frac{\partial \Pi_0}{\partial y} = \frac{2\mu}{3}(-1+2y) - 2\mu\,(y-\theta)$. Accounting for corner solutions, one immediately verifies that the agent indeed chooses the equilibrium action $\left\{x^C(\theta), y^C(\theta)\right\}$.

Next, we verify that the above contribution schedules satisfy the best response property and thus represent an equilibrium of the common influence game. By symmetry, we can focus on principal 1. Given principal 2's schedule, principal 1 is faced with a monopolistic screening problem where the agent has an *indirect* payoff function given by:

$$\begin{aligned} \Pi_{02}(x,y,\theta) \;&=\; U_0(x,y,\theta) + \alpha_0^C + \Gamma_2(x,y) \\ &= -\nu\,(x-\xi)^2 - \mu\,(y-\theta)^2 + \\ &\quad +\alpha_0^C - \frac{(1+3\Delta)\,\mu}{3}(1-y) + \frac{\mu}{3}(1-y)^2 + \pi x(1-y), \end{aligned}$$

It is evident that $\Pi_{02}(x,y,\theta)$ satisfies the sorting condition, just as $U_0(x,y,\theta)$ does, since the additional term $\Gamma_2(x,y)$ is independent of θ. Therefore, the only significant difference from the case of monopolistic influence is the participation constraint, which now is type-dependent.

To see why, note that the agent's reservation payoff is what he would obtain by accepting only principal 2's contribution and making the optimal decisions accordingly, that is, $\Pi_{02}\left(x^{EC_2}(\theta), y^{EC_2}(\theta), \theta\right)$, where:

$$\left\{x^{EC_2}(\theta), y^{EC_2}(\theta)\right\} = \arg\max_{x,y} \Pi_{02}(x,y,\theta)$$

denote the agent's optimal actions when accepting only the contribution from principal 2, given principal 2's common-influence contribution schedule. (The superscript EC_2 refers to this hypothetical scenario of *exclusivity* under the *common* influence contribution schedule.) Focusing on the case of interior solutions (the case of corner solutions can be dealt with similarly), we have

$$x^{EC_2}(\theta) = \frac{\mu(\pi(3\Delta + 6\theta - 5) - 8\nu\xi)}{3\pi^2 - 8\mu\nu},$$

and

$$y^{EC_2}(\theta) = \frac{3\pi(2\nu\xi + \pi) - 2\mu\nu(3\Delta + 6\theta - 1)}{3\pi^2 - 8\mu\nu},$$

whence:

$$\Pi_0^{EC_2}(\theta) = \alpha_0^C + \frac{1}{9\pi^2 - 24\mu\nu}\Big[\mu^2\nu\left(-\left(9\Delta^2 + 6\Delta(6\theta - 5) + 12(\theta - 1)\theta + 1\right)\right) + 6\pi\mu\nu\xi(3\Delta + 6\theta - 5) - 9\pi^2\left((\theta - 1)^2\mu + \nu\xi^2\right)\Big]$$

To address the issues raised by the type-dependent participation constraint $\Pi_0(\theta) \geq \Pi_0^{EC_2}(\theta)$, we follow an intuitive procedure that can be formally justified using the results of Jullien (2000). The procedure involves conjecturing that the solution to principal 1's problem satisfies inequality $y(\theta) \geq y^{EC_2}(\theta)$ and then verifying this conjecture ex-post. The conjecture is intuitive, as it simply states that the value of $y(\theta)$ when principal 1 is active is at least as large as when principal 1 is not active; otherwise, principal 1 would have no incentive to influence y.

To understand the implications of the conjecture, note that by the Envelope Theorem we have:

$$\frac{d\Pi_0(\theta)}{d\theta} = 2\mu\left[y(\theta) - \theta\right]$$

and

$$\frac{d\Pi_0^{EC_2}(\theta)}{d\theta} = 2\mu\left[y^{EC_2}(\theta) - \theta\right].$$

This means that inequality $y(\theta) \geq y^{EC_2}(\theta)$ is equivalent to:

$$\frac{d\Pi_0(\theta)}{d\theta} \geq \frac{d\Pi_0^{EC_2}(\theta)}{d\theta}.$$

In other words, the agent's equilibrium payoff $\Pi_0(\theta)$ must increase with θ at least as much as his reservation payoff $\Pi_0^{EC_2}(\theta)$. As a consequence, if the functions $\Pi_0(\theta)$ and $\Pi_0^{EC_2}(\theta)$ intersect, then $\Pi_0(\theta)$ must cut $\Pi_0^{EC_2}(\theta)$ from below as θ increases. Therefore, the conjecture implies that the participation constraint binds only for a single low-θ type, just as in the case where the constraint is not type-dependent.

This property, combined with the sorting condition, implies that principal 1's program reduces to pointwise maximization of the following indirect virtual surplus:

$$\begin{aligned} V_1(x, y, \theta) &= \Pi_{02}(x, y, \theta) + U_1(x, y) - \frac{1 - F(\theta)}{f(\theta)}\frac{\partial \Pi_{02}(x, y, \theta)}{\partial \theta} \\ &= -\nu(x - \xi)^2 + \pi x - \frac{1}{3}\mu\left(3\Delta + 9\theta^2 + 2y^2 - 3\theta(\Delta + 4y + 1) + 4y\right) \end{aligned}$$

The indirect virtual surplus is concave in x and y and its first derivatives are:

$$\begin{aligned}\frac{\partial V_1}{\partial x} &= 2\nu(\xi - x) + \pi \\ \frac{\partial V_1}{\partial y} &= -\frac{4}{3}\mu(-3\theta + y + 1)\end{aligned}$$

From these, we see immediately that the solution to principal 1's program entails exactly the equilibrium choices (12)-(13). At this point, it is straightforward to verify that $y(\theta) \geq y^{EC_2}(\theta)$, confirming that our solution procedure is correct.

Since the payment schedule that implements the optimal allocation for principal 1 is (27), the schedule (27) is a best response to principal 2's schedule. It is important to note that in determining the best response, we did not impose any restrictions on the set of feasible schedules for principal 1. By symmetry, the same holds true for principal 2, confirming that the payment schedules (27)-(28) and the actions (12)-(13) constitute an equilibrium of the influence game in the strong competition case.

The proof for the case of weak competition follows a similar approach. Interested readers can find details for this case in the mathematical Appendix. ■

D.3 Proof of Proposition 3 (Unilateral Incentives)

We show that starting from the common influence equilibrium of Proposition 2, under transparency, each principal has a unilateral incentive to deviate. To fix ideas, we consider the case $\pi > \frac{2\sqrt{\mu}\sqrt{\nu}}{\sqrt{3}}$, so that competition is strong, and $\Delta \geq \frac{1}{3}$. Similar arguments apply to the case where $\pi < \frac{2\sqrt{\mu}\sqrt{\nu}}{\sqrt{3}}$ or $\Delta < \frac{1}{3}$.

Let the deviating principal be principal 1. The deviation we consider consists of offering the monopolistic influence schedule, with an appropriately chosen constant payment α_0^{dev}, and imposing an exclusivity requirement whereby the agent must choose between accepting a positive payment from principal 1 or from principal 2.

With this deviation, and given that principal 2 continues to offer its common agency equilibrium schedule $\alpha_0^C + \Gamma_2^{CI}(x, y)$, the agent can choose between two options: either to accept only the contribution of principal 1, or that of principal 2. Those agent types who accept only principal 2's contribution will set $x = x^{EC_2}(\theta)$ and $y = y^{EC_2}(\theta)$, obtaining a payoff of $\Pi_0 = \Pi_0^{EC_2}(\theta)$, as detailed in the proof of Proposition 2. On the other hand, agent types who accept only the contribution of principal 1 will set $x = x^{M_1}(\theta)$ and $y = y^{M_1}(\theta)$, obtaining a payoff of:

$$\Pi_0^{M_1}(\theta) = \alpha_0^{\text{dev}} + \frac{\mu^2\nu\left[(1+\Delta)^2 - 8(1+\Delta)\theta + 8\theta^2\right] - 2\pi\mu\nu\xi(1+\Delta-4\theta) + \pi^2\left(2\theta^2\mu + \nu\xi^2\right)}{2\left(4\mu\nu - \pi^2\right)}$$

It can be easily verified that $\Pi_0^{M_1}(\theta)$ increases with θ more strongly than $\Pi_0^{EC_2}(\theta)$, implying that low-θ types will accept only the contribution of principal 2 and high-θ types only that of principal 1. The threshold between these two choices, denoted as θ^{dev}, is a decreasing function of α_0^{dev}.

At this point, it is possible to compute principal 1's deviation payoff:

$$\Pi_1^{\text{dev}} = \int_{\theta_{\min}}^{\theta^{\text{dev}}} x^{EC_2}(\theta) y^{EC_2}(\theta) f(\theta) d\theta + \\ + \int_{\theta^{\text{dev}}}^{\theta_{\max}} \left[x^{M_1}(\theta) y^{M_1}(\theta) - \Gamma^{M_1}(x^{M_1}(\theta), y^{M_1}(\theta)) - \alpha_0^{\text{dev}} \right] f(\theta) d\theta$$

and compare it to the equilibrium payoff, which is:

$$\Pi_1 = \frac{\mu \left(108\Delta^2 - 9\Delta + 2\right)}{216\Delta}.$$

The deviation payoff depends on α_0^{dev} both directly and indirectly, through θ^{dev}. To demonstrate the existence of a profitable deviation, it suffices to consider a specific value of α_0^{dev}. The mathematical Appendix focuses on the particular value of α_0^{dev} that ensures only those types for whom $y^{EC_2}(\theta) = 0$ accept principal 2's contribution. In this case, the first integral in the definition of Π_1^{dev} vanishes. Although this deviation may be suboptimal, it is always profitable, as shown through straightforward but lengthy algebra provided in the mathematical Appendix. ■

D.4 Proof of Proposition 4 (Exclusive Influence)

To show that the exclusive-influence schedules $\Gamma_i^E(x, y) = \Gamma_i^{M_i}(x, y)$ constitute an equilibrium, it suffices to note that when either principal imposes an exclusivity requirement, the agent can choose between two options: either to accept only the contribution of principal 1, or that of principal 2. Each principal must then offer the monopolistic-influence schedule to those agents who will end up in its sphere of exclusive influence. Agents who choose to accept the payment from principal 1 will then obtain a payoff of:

$$\Pi_0^{E_1}(\theta) = \alpha_{01}^E + \frac{\mu^2\nu \left[(1+\Delta)^2 - 8(1+\Delta)\theta + 8\theta^2\right] - 2\pi\mu\nu\xi(1+\Delta-4\theta) + \pi^2\left(2\theta^2\mu + \nu\xi^2\right)}{2\left(4\mu\nu - \pi^2\right)}.$$

A similar expression, with $1-\theta$ replacing θ, holds for those who choose to be influenced by principal 2. It is immediate to verify that high-θ types will opt for the contribution of principal 1, while low-θ types will opt for principal 2. The critical type $\hat{\theta}$ who is just indifferent between the two options is:

$$\hat{\theta} = \frac{1}{2} - \frac{\left(4\mu\nu - \pi^2\right)}{2\mu\left[\pi(4\nu\xi + \pi) - 4\Delta\mu\nu\right]}(\alpha_{01}^E - \alpha_{02}^E)$$

Clearly, the critical type depends on the fixed components of the contribution schedules, α_{0i}^E. (Note that the denominator of the second fraction on the right-hand side must be positive; otherwise, principals would not compete for influence.) Principal 1's payoff then becomes:

$$\Pi_1^E = \int_{\theta_{\min}}^{\hat{\theta}} \pi x^{M_2}(\theta) y^{M_2}(\theta) f(\theta) d\theta + \\ + \int_{\hat{\theta}}^{\theta_{\max}} \left[\pi x^{M_1}(\theta) y^{M_1}(\theta) - \Gamma^{M_1}(x^{M_1}(\theta), y^{M_1}(\theta)) - \alpha_{01}^E\right] f(\theta) d\theta.$$

The fixed components of the payment schedules are then determined by solving the system of first-order conditions:

$$\frac{d\Pi_i^E}{d\alpha_{0i}^E} = 0.$$

The solution is reported in the mathematical Appendix. ■

D.5 Proof of Proposition 5 (Prevalent Influence)

The proof proceeds as in the proof of Proposition 2. First, we demonstrate that, given the contribution schedules (24)-(25), the agent's optimal choices are the equilibrium actions $\{x^P(\theta), y^P(\theta)\}$. We then examine the screening problem faced by principal 1, given the equilibrium schedule of principal 2. We solve this problem by focusing on direct mechanisms and demonstrate that the solution yields the equilibrium actions. Since these actions are implemented by the equilibrium schedules, this implies that principal 1's schedule is a best response to principal 2's schedule. By symmetry, the same applies to principal 2, confirming the equilibrium under mandatory disclosure.

Consider first the agent's optimal choices given the equilibrium schedules (24)-(25). The agent must choose which principal to patronize, meaning whose sphere of prevalent influence to fall in, and what actions $x(\theta)$ and $y(\theta)$ to take within each sphere of influence. Consider the latter choice first. If the agent chooses to patronize principal 1 and thus accepts the payments $\Gamma_1^H(x, y)$ from principal 1 and $\Gamma_2^L(x, y)$ from principal 2, he maximizes:

$$\begin{aligned}
\Pi_0^{P_1}(x, y, \theta) &= U_0(x, y, \theta) + \Gamma_1^H(x, y) + \alpha_{01}^H + \Gamma_2^L(x, y) + \alpha_{02}^L \\
&= \alpha_{01}^H + \alpha_{02}^L - \nu(x - \xi)^2 - \mu(y - \theta)^2 \\
&- \frac{1}{2}\left[(1 + \Delta)\mu + \pi\xi\right] y + \frac{\left(4\mu\nu - \pi^2\right)}{8\nu} y^2 + \pi x y \\
&+ \kappa \left[\frac{\pi^2}{2\nu}(1 - y)^2 + 2\nu x^2 + 2\pi x(1 - y)\right]
\end{aligned}$$

where

$$\kappa = \frac{\sqrt{\bar{T}\upsilon}}{(2\xi\nu + \pi) + 2\sqrt{\bar{T}\nu}},$$

under the constraint:

$$\Gamma_2^L(x, y) + \alpha_{02}^L \leq \bar{T}.$$

We can focus on an unconstrained problem because the equilibrium schedules $\Gamma_i^L(x, y)$ are constructed in such a way that, when the agent maximizes the objective function without considering the constraint, the resulting solution automatically satisfies it. The partial derivatives of the objective functions are:

$$\frac{\partial \Pi_0^{P_1}(x, y, \theta)}{\partial x} = -2\nu(x - \xi) + \pi y + 4\nu\kappa x$$

and

$$\frac{\partial \Pi_0^{P_1}(x, y, \theta)}{\partial y} = -2\mu(y - \theta) - \frac{1}{2}[(1 + \Delta)\mu + \pi\xi)] + 2\frac{(4\mu\nu - \pi^2)}{8\nu}y + \pi x + \kappa\left[\frac{\pi^2}{\nu}(1 - y) - 2\pi x\right]$$

It is straightforward to verify that the optimal choices are $x(\theta) = x^P(\theta)$ and $y(\theta) = y^{M_1}(\theta)$, and that the constraint $\Gamma_2^L(x, y) + \alpha_{02}^L \leq \bar{T}$ is always satisfied as an equality. The associated payoff, $\Pi_0^{P_1}(\theta) \equiv \Pi_0^{P_1}\left(x^P(\theta), y^{M_1}(\theta), \theta\right)$, increases with θ.

Similarly, when the agent chooses to patronize principal 2, and hence chooses x and y to maximize

$$\Pi_0^{P_2}(x, y, \theta) = U_0(x, y, \theta) + \Gamma_2^H(x, y) + \alpha_{02}^H + \Gamma_1^L(x, y) + \alpha_{01}^L,$$

the outcome is $x(\theta) = x^P(\theta)$ and $y(\theta) = y^{M_2}(\theta)$, with a payoff $\Pi_0^{P_2}(\theta)$ that decreases with θ.

Since the functions $\Pi_0^{P_1}(\theta)$ and $\Pi_0^{P_2}(\theta)$ are identical upon replacing θ with $1 - \theta$, it follows that $\Pi_0^{P_1}(\theta) > \Pi_0^{P_2}(\theta)$ if and only if $\theta > \frac{1}{2}$. Thus, the agent patronizes principal 2 if $\theta < \frac{1}{2}$ and principal 1 if $\theta > \frac{1}{2}$, while being indifferent if $\theta = \frac{1}{2}$. This confirms that given the equilibrium schedules (24)-(25), the agent indeed chooses the equilibrium actions.

Next, consider principal 1's screening problem when principal 2 offers its equilibrium schedule. This is a monopolistic screening problem with an agent's indirect payoff function given by $\Pi_{02}(x, y, \theta) = U_0(x, y, \theta) + \Gamma_2(x, y) + \alpha_{02}$. This indirect payoff now takes on two values:

$$\begin{aligned}\Pi_{02}^{P_2}(x, y, \theta) &= U_0(x, y, \theta) + \Gamma_2^H(x, y) + \alpha_{02}^H \\ &= U_0(x, y, \theta) + \alpha_{02}^H - \frac{\mu}{2}(1 + \Delta + \xi)(1 - y) + \frac{4\mu^2 - \pi^2}{8\mu}(1 - y)^2 + \pi x(1 - y)\end{aligned}$$

if $T_1(\theta) \leq \bar{T}$, and:

$$\begin{aligned}\Pi_{02}^{P_1}(x, y, \theta) &= U_0(x, y, \theta) + \Gamma_2^L(x, y) + \alpha_{02}^L \\ &= U_0(x, y, \theta) + \alpha_{02}^L + \kappa\left[\frac{\pi^2}{2\nu}(1 - y)^2 + 2\nu x^2 + 2\pi x(1 - y)\right]\end{aligned}$$

where

$$\kappa = \frac{\sqrt{\bar{T}\upsilon}}{(2\xi\upsilon + \pi) + 2\sqrt{\bar{T}\upsilon}}$$

if $T_1(\theta) > \bar{T}$.

The indirect payoff function satisfies the sorting condition because $U_0(x, y, \theta)$ satisfies it and the additional terms do not depend on θ. This allows us to formulate this monopolistic screening problem as a problem of optimal control using the identity:

$$\Pi_{02}(\theta) \equiv U_{02}\left(x(\theta), y(\theta), \theta\right) + T_1(\theta),$$

which defines the agent's payoff $\Pi_{02}(\theta)$. Principal 1's payoff becomes:

$$\Pi_1 = \int_{\theta_{\min}}^{\theta_{\max}} \left[\pi x(\theta) y(\theta) + \Pi_{02}\left(x(\theta), y(\theta), \theta\right) - \Pi_{02}(\theta)\right] f(\theta)\, d\theta,$$

and the local incentive-compatibility constraints, which given the sorting condition imply the global constraints, become:

$$\frac{d\Pi_{02}(\theta)}{d\theta} = \frac{\partial \Pi_{02}(x, y, \theta)}{\partial \theta} = 2\mu\left(y - \theta\right).$$

Principal 1's problem is therefore an optimal control problem, with $\Pi_{02}(\theta)$ as the state variable and $x(\theta)$ and $y(\theta)$ as the control variables. Unlike in standard screening problems, however, this problem involves two different control systems, $(\mathcal{P}_1)$ and $(\mathcal{P}_2)$, depending on whether $\Pi_{02}(x, y, \theta) = \Pi_{02}^{P_1}(x, y, \theta)$ or $\Pi_{02}(x, y, \theta) = \Pi_{02}^{P_2}(x, y, \theta)$, and the possibility of switching from one system to the other. Technically speaking, therefore, it is a *multi-stage* optimal control problem. To solve this problem, one needs to choose a sequence of control systems, the switching points, and the control functions $x(\theta)$ and $y(\theta)$ for each system that maximize the objective function.

Besides having a different indirect payoff function, the two systems involve different participation constraints. When $\Pi_{02}(x, y, \theta) = \Pi_{02}^{P_1}(x, y, \theta)$, the agent's optimal actions when dealing only with principal 2 are:

$$\begin{aligned} x^{E_2P_1}(\theta) &= \frac{\pi k(2(\theta-1)\mu + \pi\xi) - 2\mu\nu\xi}{k\left(4\mu\nu + \pi^2\right) - 2\mu\nu} \\ y^{E_2P_1}(\theta) &= \frac{k\left(4\theta\mu\nu + 2\pi\nu\xi + \pi^2\right) - 2\theta\mu\nu}{k\left(4\mu\nu + \pi^2\right) - 2\mu\nu}, \end{aligned}$$

where the superscript E_2P_1 stands for the case where the agent takes only the payment from principal 2 intended for agents who fall in the sphere of prevalent influence of principal 1, and the agent's reservation payoff is:

$$\Pi_{02}^{E_2P_1}(\theta) \equiv \Pi_{02}\left(x^{E_2P_1}(\theta), y^{E_2P_1}(\theta), \theta\right).$$

Similarly, when $\Pi_{02}(x, y, \theta) = \Pi_{02}^{P_2}(x, y, \theta)$, the agent's optimal actions when accepting only principal 2's payments are $x^{E_2P_2}(\theta) = x^{M_2}(\theta)$ and $y^{E_2P_2}(\theta) = y^{M_2}(\theta)$, so and the agent's reservation payoff is:

$$\Pi_0^{E_2P_2}(\theta) \equiv \Pi_{02}\left(x^{M_2}(\theta), y^{M_2}(\theta)\right).$$

Therefore, the participation constraint is $\Pi_{02}(\theta) \geq \Pi_0^{E_2P_1}(\theta)$ in system $(\mathcal{P}_1)$ and $\Pi_{02}(\theta) \geq \Pi_0^{E_2P_2}(\theta)$ in system $(\mathcal{P}_2)$.

Furthermore, while problem $(\mathcal{P}_1)$ includes only the standard participation and incentive-compatibility constraints, problem $(\mathcal{P}_2)$ has an additional constraint on the level of payments, $T_1(\theta) \leq \bar{T}$. This constraint can be rewritten as

$$\Pi_{02}(\theta) \leq \bar{T} + \Pi_{02}^{P_2}(x, y, \theta),$$

indicating that in problem $(\mathcal{P}_2)$, there is both a lower and an upper bound on the state variable $\Pi_{02}(\theta)$.

Finally, an additional incentive compatibility constraint that must be satisfied is that the control system chosen for a type θ must guarantee to him a (weakly) greater utility than the other. This reflects the fact that the agent freely chooses whose contributions to accept.

To proceed, note that the state variable enters linearly, and with the same coefficient, in both control problems $(\mathcal{P}_1)$ and $(\mathcal{P}_2)$. This implies that at the optimum, the costate variable is independent of both the specific sequence of the control systems and the switching points. Intuitively, since the only reason why the allocation for a given type θ is distorted is to reduce the information rent for closer types, what matters is how many closer types there are relative to the current type, i.e. the hazard rate $\frac{1-F(\theta)}{f(\theta)}$ for principal 1, and $\frac{F(\theta)}{f(\theta)}$ for principal 2. The precise allocations designed for closer (and more distant) types do not matter.

This greatly simplifies finding the solution to this multi-stage problem. It implies that for any possible sequence of control systems and switching points, the optimal control function for the multi-stage problem coincides with $\left\{x^{P_1}(\theta), y^{P_1}(\theta)\right\}$ whenever system$(\mathcal{P}_1)$ applies, and with $\left\{x^{P_2}(\theta), y^{P_2}(\theta)\right\}$ whenever system $(\mathcal{P}_2)$ applies, where $\left\{x^{P_i}(\theta), y^{P_i}(\theta)\right\}$ denotes the solution that would be obtained if system $(\mathcal{P}_i)$ were applied to all types θ. In other words, the optimal control function for the multi-stage problem is constructed by appropriately combining the optimal control functions for problems $(\mathcal{P}_1)$ and $(\mathcal{P}_2)$ taken separately.

This separation property also allows us to determine the optimal sequence of control problems. This is because the slope of the agent's equilibrium payoff function $\Pi_{02}(\theta)$ is $2\mu(y(\theta) - \theta)$ by the incentive compatibility constraint, and it is straightforward to verify that $y^{P_1}(\theta) > y^{P_2}(\theta)$. That is, y is higher if principal 1 acts as primary contributor (when $y = y^{M_1}(\theta)$) rather than as secondary contributor (when $y = y^{M_2}(\theta)$). As a consequence, the rate of change of the equilibrium payoff $\Pi_{02}(\theta)$ with respect to θ is always higher under system $(\mathcal{P}_1)$ than under system $(\mathcal{P}_2)$. Since $\Pi_{02}(\theta)$ must be continuous, this implies that there will be at most one switch as θ increases, from system $(\mathcal{P}_2)$ to system $(\mathcal{P}_1)$.

Let us denote the switching point as $\hat{\theta}$. For the time being, we treat $\hat{\theta}$ as fixed and separately characterize the solution to the control problem $(\mathcal{P}_1)$, which applies for $\theta > \hat{\theta}$, and problem $(\mathcal{P}_2)$, which applies for $\theta < \hat{\theta}$. Later we will consider the optimal choice of $\hat{\theta}$.

Problem $(\mathcal{P}_1)$ is a standard monopolistic screening problem: find the direct mechanism

$(x(\theta), y(\theta))$ that solves:

$$\max \int_{\theta_{\min}}^{\hat{\theta}} \Pi_1^{\mathcal{P}_1}(x(\theta), y(\theta), \theta) f(\theta)\, d\theta$$

where $\Pi_1^{\mathcal{P}_1}(x(\theta), y(\theta), \theta) \equiv \pi x(\theta) y(\theta) + \Pi_{02}^{P_1}(x(\theta), y(\theta), \theta) - \Pi_{02}(\theta)$ is the payoff principal 1 obtains from type θ, subject to the standard incentive-compatibility and participation constraints:

$$\begin{aligned} \frac{d\Pi_{02}(\theta)}{d\theta} &= 2\mu\,(y - \theta) \\ \Pi_0(\theta) &\geq \Pi_0^{E_2 P_1}(\theta). \end{aligned}$$

As in the proof of Proposition 2, we conjecture that the participation constraint binds only at the bottom of the distribution of types to whom the problem applies, $\theta = \hat{\theta}$, and verify the conjecture afterward. By standard arguments, it then follows that the problem reduces to the pointwise maximization of the virtual surplus:

$$V_{02}^{P_1}(x, y, \theta) = \Pi_{02}^{P_1}(x, y, \theta) + U_1(x, y) - \frac{1 - F(\theta)}{f(\theta)} \frac{\partial U_0(x, y, \theta)}{\partial \theta}.$$

The virtual surplus function is concave, and from the first-order conditions, it is straightforward to verify that the optimal actions are $x(\theta) = x^P(\theta)$ and $y(\theta) = y^{M_1}(\theta)$. It is also straightforward to verify that $y^{P_1}(\theta) = y^{M_1}(\theta) > y^{E_2 P_2}(\theta)$, which confirms that the participation constraint binds only at $\theta = \hat{\theta}$. Since the optimal actions are implemented by the schedule $\Gamma_1^H(x, y)$ of the primary principal, as verified above, this schedule represents a best response to $\Gamma_2^L(x, y)$.

Consider next problem $(\mathcal{P}_2)$. The payoff function now is $\Pi_{02}^{P_2}(x, y, \theta)$, and the participation constraint is $\Pi_{02}(\theta) \geq \Pi_{02}^{E_2 P_2}(\theta)$, but now we also have the additional constraint on the level of payments, $T_1(\theta) \leq \bar{T}$ or, equivalently, $\Pi_0(\theta) = T_1(\theta) + \Pi_{02}^{P_2}(x, y, \theta) \leq \bar{T} + \Pi_{02}^{P_2}(x, y, \theta)$. Thus, problem $(\mathcal{P}_2)$ can be summarized as follows: find the direct mechanism $(x(\theta), y(\theta))$ that solves:

$$\max \int_{\theta_{\min}}^{\hat{\theta}} \Pi_1^{\mathcal{P}_2}\,(x(\theta), y(\theta), \theta)\, f(\theta)\, d\theta$$

where $\Pi_1^{\mathcal{P}_2}\,(x(\theta), y(\theta), \theta) \equiv \pi x(\theta) y(\theta) + \Pi_{02}^{P_2}(x(\theta), y(\theta), \theta) - \Pi_{02}(\theta)$, subject to:

$$\begin{aligned} \frac{d\Pi_0(\theta)}{d\theta} &= 2\mu\,(y - \theta) \\ \Pi_0(\theta) &\geq \Pi_0^{E_2 P_2}(\theta) \\ \Pi_0(\theta) &\leq \bar{T} + \Pi_{02}^{P_2}(x, y, \theta). \end{aligned}$$

The first constraint is the incentive compatibility constraint, the second is the participation constraint, and the third is the constraint on the level of payments.

First of all, it is clear that in equilibrium, principal 1 will not attempt to influence the choice of y for those agents who fall within the sphere of prevalent influence of principal 2. Otherwise, principal 2 would have an incentive to lower the cap $\bar{T}_2$ on the payments that can be accepted from the secondary contributor—i.e., principal 1—thereby reducing principal 1's influence on y and increasing principal 2's payoff.

More precisely, when principal 2 lowers $\bar{T}_2$, causing the agent to decrease y, this reduces the agent's payoff within the sphere of principal 2's influence, potentially decreasing the switching type $\hat{\theta}$. However, since the agent optimizes the choice of y, a small decrease in y has only a second-order effect on the agent's payoff, and thus a second-order effect on $\hat{\theta}$. By contrast, the reduction in y has a first-order positive effect on principal 2's payoff, which would therefore increase.

This implies that, within the sphere of prevalent influence of principal 2, we must have $y^P(\theta) = y^{M_2}(\theta)$. This has two important consequences. First, it implies that the participation constraint must be binding for all types. This is because when $y^P(\theta) = y^{M_2}(\theta)$, the slope of the agent's equilibrium payoff $\Pi_{02}(\theta)$—which is $y^P(\theta)$—coincides with the slope of the reservation payoff, $y^{M_2}(\theta)$.

Second, the constraint on the level of payments, $T_1(\theta) \leq \bar{T}$, must also be binding for all types. To see why, suppose instead that the payment constraint is slack for some types. Then these types would influence the choice of x as in the common-influence equilibrium, so that $x(\theta) = x^C(\theta)$. However, recall that $x^C(\theta)$ is in fact independent of θ, implying that the rate of change of $\Pi_{02}^{P_2}(x, y, \theta)$ must match the rate of change of $\Pi_{02}(\theta)$. Since the payment constraint is written as $\Pi_0(\theta) \leq \bar{T} + \Pi_{02}^{P_2}(x, y, \theta)$, this implies that the constraint must either be always slack (which is evidently impossible) or always binding.

Therefore, the solution can be obtained directly from the condition $\bar{T} + \Pi_{02}^{P_2}(x, y^{M_2}, \theta) = \Pi_0^{E_2P_2}(\theta)$, which yields the actions $\{x^P(\theta), y^P(\theta)\}$. Since these actions are implemented by the prevalent-influence equilibrium schedule of the secondary principal, this schedule represents a best response to that of the primary principal.

Let us now consider the optimal choice of $\hat{\theta}$. At the optimum, the following conditions must hold:

$$\Pi_{02}^{P_2}(\hat{\theta}) = \Pi_{02}^{P_1}(\hat{\theta}) \tag{29}$$

and:

$$\frac{\Pi_1^{\mathcal{P}_1}\left(x^{M_1}(\hat{\theta}) + \sqrt{\frac{\bar{T}}{\mu}}, y^{M_1}(\hat{\theta}), \hat{\theta}\right) - \Pi_1^{\mathcal{P}_2}\left(x^{M_2}(\hat{\theta}) + \sqrt{\frac{\bar{T}}{\mu}}, y^{M_2}(\hat{\theta}), \hat{\theta}\right)}{2\mu\left(y^{M_1}(\hat{\theta}) - y^{M_2}(\hat{\theta})\right)} = \frac{1 - F(\hat{\theta})}{f(\hat{\theta})}.$$

The first condition, which states that the agent must be indifferent between falling within the sphere of prevalent influence of principal 1 or principal 2, captures the new incentive compatibility constraint mentioned earlier regarding the choice of the control system. The second condition requires the continuity of the Hamiltonian, a standard necessary condition in multi-stage optimal control theory, often seen as a generalization of the Weierstrass-Erdmann corner conditions. It is derived from the condition that principal 1 optimally chooses the constant α_{01}^H. Any increase in this constant term applies to a mass $1 - F(\hat{\theta})$ of agents, decreasing the principal payoff, but on the other hand it lowers the critical threshold $\hat{\theta}$, implying that principal 1 gains the difference $\Pi_1^{\mathcal{P}_1}(\hat{\theta}) - \Pi_1^{\mathcal{P}_2}(\hat{\theta})$ on a mass $f(\hat{\theta})\frac{d\hat{\theta}}{d\alpha_0^H}$. From condition (29), one sees that the derivative is $\frac{d\hat{\theta}}{d\alpha_0^H} = \frac{1}{2\mu\left(y^{M_1}(\hat{\theta}) - y^{M_2}(\hat{\theta})\right)}$.

The Weierstrass-Erdmann corner condition implies that marginal profitability always increases at a switching point. The economic intuition is straightforward: consider an increase in the constant term of the tariff $\Gamma_1^H(x, y)$, which applies when $\theta > \hat{\theta}$. (Notice that if the entire schedule $\Gamma_1^H(x, y)$ is shifted upward by a constant, the equilibrium choices within that interval remain unchanged.) This shift has a direct negative impact on profits extracted from higher types, as well as an indirect effect due to the resulting decrease in $\hat{\theta}$. The indirect effect would disappear if the numerator on the left-hand side were zero. However, at the optimum, the indirect effect must be positive to precisely offset the negative direct effect. This implies that profitability must be higher to the right than to the left of the switching point. Therefore, since the denominator on the left-hand side is positive, the system optimally switches from (P_2) to (P_1), with principal 1 gaining more at the margin if type $\hat{\theta}$ falls within its sphere of influence rather than that of principal 2.

The equilibrium values of $\hat{\theta}$ and the constants α_{0i}^H are determined by the above conditions for principal 1 and the analog conditions for principal 2. Clearly, in the symmetric case it must be $\hat{\theta} = \frac{1}{2}$. On the other hand, the constants α_{0i}^H must be determined simultaneously with the constants α_{0i}^L. The latter constant are pinned down by the participation constraint in the screening problem faced by principals when they act as secondary contributors. The exact formulas are reported in the mathematical Appendix.

It remains to prove the claim that $\bar{T}$ must belong to the interval $\left[0, \hat{T}\right]$, and that there are no equilibria when $\bar{T} > \hat{T}$. First of all, note that $\hat{T}$ is determined by the condition that $x^P(\theta_{\min}) = x^P(\theta_{\max}) = x^C$. If $\bar{T} > \hat{T}$, there would be a non-empty set of extremist types who set $x = x^C$, as a higher value of x is never optimal, and obtain a payment lower than $\bar{T}$. But then the secondary principals would have room for influencing the agent's choice of y without the agent violating the constraint on the level of payments set by the primary principals. However, we have argued above that this is inconsistent with equilibrium. Therefore, the equilibrium requires that $\bar{T} \in \left[0, \hat{T}\right]$. Note that for certain parameter values, the interval $\left[0, \hat{T}\right]$ collapses, in which case only the exclusive-influence equilibrium exists.

This completes the proof of Proposition 5. ■

E Three principals

Extending the model to $n \geq 3$ symmetric principals is not straightforward, as it introduces multi-dimensional uncertainty about the doctor's ideal market shares across the different brands. To preserve tractability and retain one-dimensional uncertainty, symmetry must be relaxed. For example, one can assume the existence of three drugs, but posit that the doctor's ideal market share for drug 3 is common knowledge, while uncertainty applies only to his preferences over drugs 1 and 2. As discussed in the main text, this may offer a realistic approximation of the NOACs market during the period under study.

Consider, then, a variant of the model where, after selecting x, representing the proportion of patients receiving branded drugs, the doctor further chooses two variables, z and y, where z denotes the proportion of treated patients who take drug 3, while y denotes the share of the remaining patients, that is, $x(1-z)$, who receive drug 1. The doctor's intrinsic utility now is:

$$U_0 = -\nu\left(x-\xi\right)^2 - \eta\left(z-\zeta\right)^2 - \mu\left(y-\theta\right)^2,$$

where ζ is the doctor's ideal share of drug 3, which is common knowledge. The pharmaceutical companies' gross profits become:

$$\begin{aligned} U_1 &= \pi x(1-z)y \\ U_2 &= \pi x(1-z)(1-y) \\ U_3 &= \pi x z. \end{aligned}$$

Since there is no uncertainty along the z-dimension, the optimal contribution schedule for firm 3 is the truthful one:

$$T_3 = \pi x z + \text{ constant},$$

both with and without transparency. One can then define an indirect utility function

$$\begin{aligned} \tilde{U}_0 &= \max_z\left[U_0 - T_3\right] \\ &= -\nu\left(x-\xi\right)^2 - \mu\left(y-\theta\right)^2 + \frac{\pi x\left(\pi x + 4\zeta\eta\right)}{4\eta} + \text{ constant}, \end{aligned} \tag{30}$$

and analyze the competition for influence between principal 1 and 2 proceeding as in the baseline model. Since the extra-term in (xx) does not involve y, the algebra becomes more cumbersome, but all the qualitative results derived in the baseline model extend to this new setting.

F Descriptive statistics

In this appendix, we present some descriptive statistics.

Table 8, Panel A presents various relevant characteristics of the cardiologists included in our study. Only 10% of the doctors in our sample are female,[46] which is lower than the overall share of female physicians in the profession (50%).[47]

Furthermore, 96% of the doctors in our sample hold an MD degree as opposed to an osteopathic physician degree (OS), and their average year of graduation is approximately

[46] This is why we use masculine pronouns for the agents in this paper.

[47] This gender disparity is notable compared to other specialties. However, the average graduation year for female doctors in our sample is 1995, while the overall sample graduated in 1988. This suggests that the share of female cardiologists is likely to increase in the future.

1988. This means that the average doctor has around 25 years of experience at the start of our study period.

Table 8: Descriptive statistics

Panel A: Characteristics of the physicians, N=16674				
	Mean	Std. dev.	Min	Max
Share of female	0.10	-	-	-
Share of medical doctors	0.96	-	-	-
Graduation year	1987.83	10.32	1955	2014
Panel B: Prescriptions made by physicians (yearly)				
	Mean	Std. dev.	Min	Max
Amount prescribed - firm level	27633.65	46295.46	0.00	1.8e+06
Total amount prescribed	1.1e+05	1.1e+05	5002.93	3.0e+06
Number of anticoagulants prescribed	3.31	0.81	1.00	4.00
- Share 0 anticoagulant	0.00	-	-	-
- Share 1 anticoagulant	0.03	-	-	-
- Share 2 anticoagulants	0.15	-	-	-
- Share 3 anticoagulants	0.32	-	-	-
- Share 4 anticoagulants	0.51	-	-	-
HHI prescriptions	5282	1767	2520	10000
Panel C: Payments received by physicians (yearly)				
	Mean	Std. dev.	Min	Max
Amount (doctor-firm)	236.84	2779.53	0.00	3.0e+05
Amount (doctor-firm, and positive)	683.16	4688.25	0.06	3.0e+05
Total amount received	947.37	6043.14	0.00	3.0e+05
Total amount received (if positive)	1502.03	7554.29	0.97	3.0e+05
Number of firms making payments	1.39	1.24	0.00	3.00
- Share 0 payers	0.37	-	-	-
- Share 1 payers	0.16	-	-	-
- Share 2 payers	0.19	-	-	-
- Share 3 payers	0.28	-	-	-
HHI payments	4165	3726	0	10000
HHI payments (at least 1 payer)	6603	2431	3333	10000

Notes: *Share of medical doctors* is the share of doctors with an MD degree among those doctors for which this information is reported. The share of doctors in *Physician Compare* without this information 73.4%. In our data it is lower but still sizable: 47.8%. The *Graduation year* is instead reported for the vast majority of the physicians (approximately 99.2%). Approx. 13% graduated from a top-20 medical school.

In the study period, cardiologists had a choice of four different anticoagulants: *Pradaxa*, *Xarelto*, *Eliquis*, and *Warfarin* (generic). Panel B of Table 8 provides summary statistics on prescriptions. On average, each doctor prescribed 3.31 different anticoagulants per year, for a total value of $110,534. The table also shows the percentage of doctors who prescribed only one product (3%), two (15%), three (32%), or all four (51%) in a year. The HHI index of prescriptions (scaled between 0 and 10,000) is also reported in this panel.

Panel C of Table 8 provides an overview of payments received by cardiologists in our sample. As anticipated, generic manufacturers do not make payments to doctors. A significant proportion of physicians (37%) also do not receive any payments from branded drug companies. On average, doctors receive $947 per year related to the promotion of anticoagulants. Among those who receive such payments, the average yearly amount is $1,502.

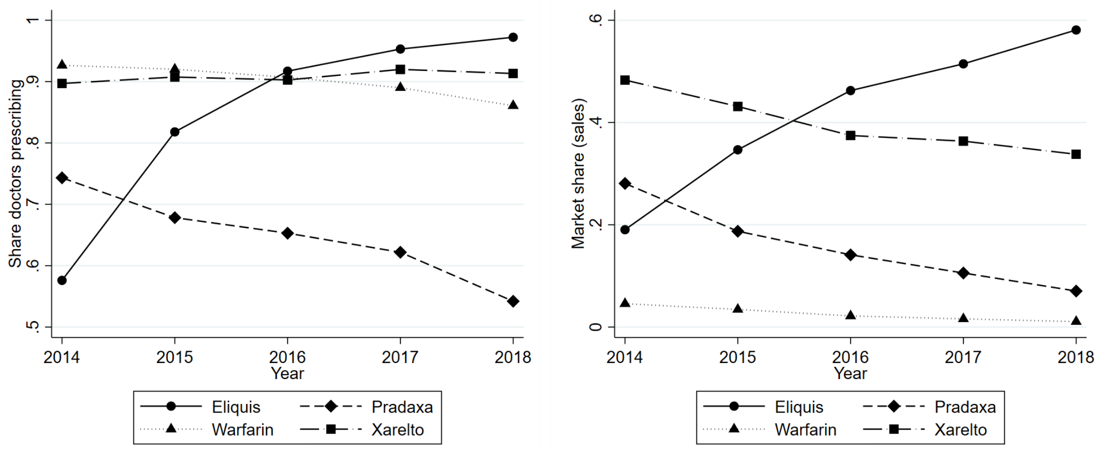

Figure 6: Left panel: Share of doctors making at least one prescription for each anticoagulant. Right panel: market shares (in value) of the four anticoagulants from 2014 to 2018.

Figure 6 presents some of the trends discussed above. The left panel shows the percentage of doctors who prescribe a positive amount of each anticoagulant in any given year. According to the graph, in 2018, *Warfarin* was still prescribed at least once in a year by more than 80% of doctors, down from over 90% in 2014 and 2015. Among the branded drugs, *Xarelto* was the most widely prescribed at the beginning of the study period but was later surpassed by *Eliquis*. Over time, *Pradaxa* continued to lose popularity, and in 2018, it was prescribed by slighty more than 50% of doctors.

The right panel of Figure 6 shows the market shares of the drugs in terms of sales value. *Warfarin*'s market share is negligible, due to its low price compared to the branded products. *Eliquis*, *Xarelto*, and *Pradaxa* show a similar pattern as in the left panel, with *Eliquis* gaining market share rapidly, *Xarelto* being the second-largest seller, and *Pradaxa* rapidly losing market share from almost 30% in 2014 to less than 10% in 2018.

The left panel of Figure 7 displays the share of doctors receiving payments from each of the three branded products between 2014 and 2018. Notably, the share of doctors receiving payments is decreasing over time for all three products. While the safety concerns associated with *Pradaxa* may explain the declining trend for that drug, the decreasing trend for *Eliquis* and *Xarelto* is surprising given their recent commercialization. One may wonder that transparency regulations might have played a role here.

This hypothesis could also provide an explanation for why the share of doctors receiving payments from all companies decreases even more rapidly than that of individual drugs, by

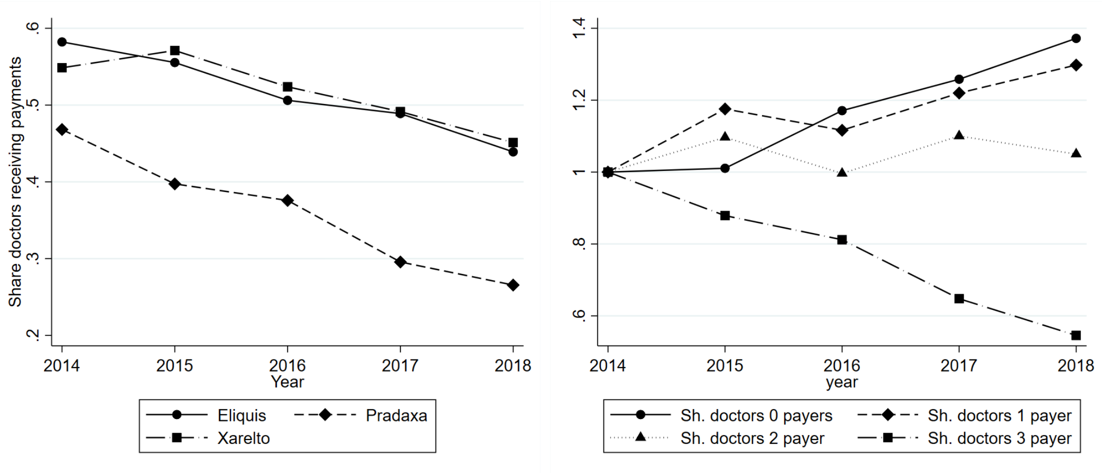

Figure 7: Left panel: share of physicians receiving payments by Eliquis, Pradaxa and Xarelto between 2014 and 2018. Right panel: share of doctors not receiving any payment, and shares of doctors receiving payments from 1, 2 or 3 companies bewteen 2014 and 2018.

around 40%. This is consistent with the polarization of payments predicted by our theory, which can also explain why the proportion of doctors receiving contributions from only one company is increasing.

G Additional Tables and Figures

In this section we collect additional tables and figures referenced in the main text of the paper. Section G.0.1 reports the analysis on payments to support the assumption of absence of pre-treatment trends.

G.0.1 Pre-trends in payments

In this section we report the analysis of payments to test the assumption of common trend between doctors in treated and control areas. We test the assumption of common trend on payment data at the doctor-month level (i.e., at a temporally more disaggregated level than in the main analysis where we are forced to aggregate at the annual level because the prescription data is provided at such level). First, we plot the average (monthly) payments received by doctors in the two groups of States and secondly, we estimate a simple model for payments to test whether the trends (if present) have different slopes.

The left panel of Figure 8 reports the average of monthly payments for the doctors living in treated states (solid line) and control states (dashed line), while the right panel of the figure shows the difference over time between these average payments made to the two groups of doctors. In sum, payments display positive trends in both groups of doctors, probably following a longer trend in marketing of the new drugs by the respective manufacturers.

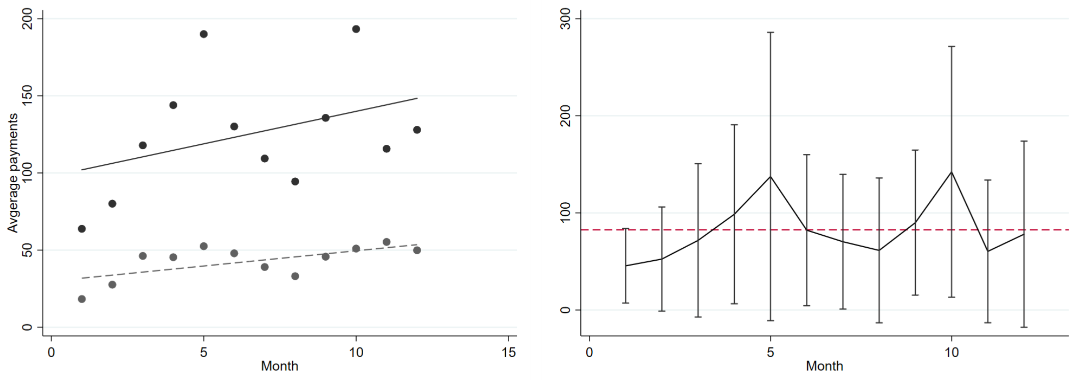

Figure 8: Left panel: Average monthly payments received by doctors in treated States (solid line) and in control States (dashed line), in the year before the introduction of transparency regulation. Right panel: Differences in monthly payments received by doctors in the treated States and in control States, in the year before the introduction of transparency regulation (vertical bars indicate 10% confidence intervals).

However, the right panel of Figure 8 also shows that the gap between the two groups does not seem to increase significantly over time, thus suggesting that the slopes of the two trends are not different.

The model that we employ to test the hypothesis of common trend is reported in equation (31):

$$Monthly\ Payments_{ist} = \gamma_1 Treated_s \times Trend_t + \gamma_0 Trend_t + \eta_i + \varepsilon_{ist} \tag{31}$$

where i identifies the doctor, s is the State where the doctor works, and t is the month. The model includes doctor fixed-effects η_i and a time-trend (monthly) $Trend_t$, thus the focus is on the estimated coefficient of the interaction term $Treated_s \times Trend_t$.

Table 9 reports the estimated coefficients of model (31) on the payments pre-transparency. Column (1) reports the estimates when we employ the full sample of doctors while column (2) reports the estimates of the model when we restrict the analysis to the subset of doctors receiving positive payments.

Results indicate that payments of both groups of doctors are characterized by a positive (upward) trend but do not follow different trends between the two groups of doctors in the 12 months before the introduction of transparency regulation. However, this analysis has one important limitations, namely the lack of a long pre-treatment period, as data is available only for 12 months before the introduction of transparency regulation. Keeping in mind such limitation, we provide evidence rejecting the presence of different trends pre-treatment, which supports our identification. Despite not being fully conclusive evidence, finding a large difference in the pattern of payments in the year before the adoption of transparency, would cast concerns on the identification strategy.

Table 9: Payments before the introduction of transparency regulation

Dep. variable: Payments	Full sample	Payments > 0
	(1)	(2)
Treated × Trend	2.239	5.842
	(3.566)	(11.390)
Trend	1.975***	6.350***
	(0.226)	(0.856)
Constant	30.823***	105.760***
	(1.485)	(5.591)
Doctor FE	YES	YES
R^2	0.202	0.388
Observations	359136	106696

Standard errors are clustered at the physician level. *, **, and *** denote significance at the 10, 5, and 1 percent level, respectively.

H Analysis of the value of prescriptions for companies

This section derives a transparent measure of the company's operating gain per treated U.S. patient-year in 2014 for two direct oral anticoagulants, apixaban (Eliquis) and rivaroxaban (Xarelto). The analysis proceeds from observable patient list spend to operating profit, with explicit assumptions on rebates and discounts, manufacturing costs, commercialization expenses, and real-world persistence.

Define the annual list price per continuously treated patient as P_L.

The fraction of list that is actually realized by the manufacturer after rebates, statutory discounts, chargebacks, and price concessions is $\alpha \in (0, 1)$, so that the manufacturer's net revenue before operating expenses equals αP_L for a full year of therapy. Brand drugs' "gross-to-net" discounts (rebates to PBMs/insurers, statutory discounts, patient support, etc.) were meaningful by 2014, though smaller than today. Broad evidence over 2007–2018 shows list prices rose much faster than net prices, implying sizable average discounts; by the mid-2010s, net prices were commonly 50–80% of list, varying by product and payer mix. Wholesaler margins on small-molecule brands are only 2–4% of net and can be ignored for a first-pass patient-level estimate.

The manufacturing cost of goods (COGS, like active substance, tableting, packaging, and quality release) for a full year is c_{full}. Because real-world exposure often falls short of a full year, we introduce a persistence factor $\gamma \in (0, 1]$ that scales both revenue and physical COGS. For example, for Eliquis, a small-molecule tablet, we have 5 mg twice daily, implying 10 mg/day and 3.65 g API/year). Export data and price trackers put apixaban API in the mid-\$3,000–\$6,500/kg range in recent years—orders of magnitude that plausibly bracket 2014 cost, too. Even at \$5,000/kg, the API per patient-year aproximately \$18 (3.65 g × \$5/g). Tableting/packaging for oral solids is often a few cents per tablet; many analyses show feasible generic tableting costs below \$0.025–\$0.10 per tablet. With 730 tablets/year, that's roughly \$18–\$73 for conversion/packaging/QC. Putting it together, COGS in \$50–\$150 per patient-year is a reasonable band.

Commercialization and in-market evidence-generation expenses (selling, general, and administrative together with post-approval RD) are modeled as a share $\beta \in (0,1)$ of net revenue, reflecting that these period costs scale with the intensity of commercialization rather than with grams of active substance.

With these primitives, net revenue, cost of goods, operating expenses, and operating profit per patient-year are

$$R = \alpha\gamma P_L, \qquad C = \gamma c_{\text{full}}, \qquad O = \beta R = \beta\alpha\gamma P_L,$$

and thus a value per patient/year

$$\pi = R - C - O = (1-\beta)\alpha\gamma P_L - \gamma c_{\text{full}}.$$

For apixaban, the standard stroke-prevention dose is 5, mg twice daily (dose reductions apply in specific subgroups; see EMA). Contemporary monthly list prices near \$287–\$289 imply $P_L \approx \$3450$ for a full-year exposure. Reasonable ranges are $\alpha \in [0.55,,0.75]$ to span 2014 gross-to-net dynamics for oral brands, $\beta \in [0.30,,0.50]$ to reflect commercialization and medical evidence burdens in an early growth phase, $\gamma \in (0,1]$ for persistence, and $c_{\text{full}} \in [50,150]$ given very low grams of Active Pharmaceutical Ingredient (API) and inexpensive oral-solid conversion. A central case with $\alpha = 0.65,\ \beta = 0.40,\ \gamma = 1,\ c_{\text{full}} = \$100$ yields

$$\Pi = (1-\beta)\alpha\gamma P_L - \gamma c_{\text{full}} = 2242.50 - 897.00 - 100.00 = \$1245.50$$

per patient-year.

For rivaroxaban, the typical NVAF dose is 20, mg once daily with food; for VTE treatment, 15, mg twice daily for 21 days then 20, mg once daily (see EMA). A 2014-level daily list price near \$10.49 implies $P_L \approx 365 \times 10.49 = \$3828.85$ for full-year exposure. We maintain $\alpha \in [0.55, 0.75]$, $\beta \in [0.30, 0.50]$, and $\gamma \in (0,1]$ as above. Because annual grams of API are roughly double apixaban's, we set $c_{\text{full}} \in [60, 180]$. A central case with $\alpha = 0.65,\ \beta = 0.40,\ \gamma = 1,\ c_{\text{full}} = \$120$ gives

$$\Pi = 2488.75 - 995.50 - 120.00 = \$1373.25.$$

We now consider empirical evidence indicating an operating gain of approximately $\pi = \$500$ per patient-year. The profit identity sets the necessary composite of net capture and persistence. For apixaban with $P_L = \$3450$ and $\gamma c_{\text{full}} \approx \$100$, one realistic 2014 constellation that attains this value is $\alpha = 0.55,\ \beta = 0.60,\ \gamma = 0.80,\ c_{\text{full}} = \$120$, which produces

$$\alpha\gamma P_L = 0.55 \times 0.80 \times 3450 = \$1518.00, \qquad \beta\alpha\gamma P_L = 0.60 \times 1518.00 = \$910.80,$$

and thus

$$\Pi = 1518.00 - 910.80 - 96.00 = \$511.20.$$

For rivaroxaban with $P_L = \$3828.85$ and $\gamma c_{\text{full}} \approx \$120$, a plausible combination is $\alpha = 0.50,\ \beta = 0.60,\ \gamma = 0.75,\ c_{\text{full}} = \$150$, yielding

$$\Pi = 1435.82 - 861.49 - 112.50 = \$461.83.$$

These calculations show that the estimated value of $\pi = \$500$ is consistent with a reasonable assessment of the different costs.