

Imperfect Information, Social Norms, and Beliefs in Networks*

Theodoros Rapanos^{†1}, Marc Sommer^{‡2}, and Yves Zenou^{§3}

¹*Department of Economics, Södertörn University*

²*University of Zurich*

³*Department of Economics, Monash University*

October 29, 2019

Abstract

We develop a simple Bayesian network game in which players, embedded in a network of social interactions, bear a cost from deviating from the social norm of their peers. All agents face uncertainty about the private benefits and the private and social costs of their actions. We prove the existence and uniqueness of a Bayesian Nash equilibrium and characterize players' optimal actions. We then show that denser networks do not necessarily increase agents' actions and welfare. We also find that, in some cases, it is optimal for the planner to affect the payoffs of selected individuals rather than all agents in the network. We finally show that having more information is not always beneficial to agents and can, in fact, reduce their welfare. We illustrate all our results in the context of criminal networks in which offenders do not know with certitude the probability of being caught and do not want to be different from their peers in terms of criminal activities.

JEL classification: C72, D82, D85, K42.

Keywords: Bayesian games, beliefs, network, conformism, value of information, crime.

*We are grateful to the Jan Wallander and Tom Hedelius foundation and the Tore Browaldh foundation for funding our work. We also thank the participants of the 17th Conference on Research on Economic Theory and Econometrics, Tinos, Crete, 2018, and the departmental seminar at Södertörn University, Sweden, for their helpful comments.

[†]*Email:* theodoros.rapanos@sh.se.

[‡]*Email:* marc.sommer@yetnet.ch.

[§]*Email:* yves.zenou@monash.edu.

1 Introduction

Social and professional networks are pervasive in people’s everyday lives and directly or indirectly influence their choices and their behavior. For example, the decision of an individual to pursue college education, consume a product, work hard or shirk, and engage in criminal activities, has been shown to be affected by their social environment. An important aspect of these decisions is that agents do not have a perfect knowledge of their environment. For example, when someone decides to commit a crime, she does know with certainty the probability of being caught and, if caught, the severity of the punishment. Students studying at universities do not exactly know the returns from education, etc. Therefore, they need to form expectations about these aspects and, in a social network context where agents are affected by their peers, they also need to form expectations about their peers’ actions, their peers’ peers’ actions, and so forth (i.e., *higher order beliefs*).

In this paper, we develop a simple Bayesian network game in which players, embedded in a network of social interactions, bear a cost from deviating from the social norm of their peers. They face uncertainty about the private benefits and the private and social costs of their actions. These parameters are potentially heterogeneous across players and consist of two components: a *global* one that is common to all players, which represents the universal or invariant payoff or cost resulting from an action, and an *idiosyncratic* one, which captures the part of benefit or cost of the action that depends on the individual characteristics of each individual. To model uncertainty, each player is assumed to observe some (potentially noisy) signal about the value of her payoff parameters, based on which she optimally chooses an action. Even though agents are only directly affected by the choices of their immediate neighbors in the network, they also need to infer the actions of players located more than one link away, since these “higher-order neighbors” may affect their immediate neighbors, who in turn affect them. Because of linear best responses, Nash equilibrium outcomes will depend on higher-order average expectations.

We first establish the existence and uniqueness of a Bayesian Nash equilibrium and characterize players’ optimal actions. We show, in particular, that what determines the equilibrium action of each player depends on the accuracy of the signal she receives (which impacts the way she forms expectations), her position in the network (which affects her beliefs about her neighbors’ actions and thus her social norm, which is the average of the actions of her neighbors) and, of course, the states of the world (i.e., the values of the different parameters of her utility function).

Second, we perform some comparative statics analyses to examine how the social environment, and the idiosyncratic characteristics of the players, affect their equilibrium actions and welfare. One interesting result is to show that an upward shift in the *idiosyncratic* component of a given player’s *social cost* (taste for conformity) parameter leads to an increase in this player’s action if her action is greater than that of her social norm. If we now consider an upward shift in the *global* component of the social cost parameter, which affects all players in the same way, then the marginal social costs increase for players whose actions are greater than that of their social norms, which leads to a reduction in equilibrium action. A similar result has been obtained in the perfect information case (Ushchev and Zenou (2019)). However, if we assume that signals and states are stochastically independent, then we can show that the impact of a First-Order Stochastic Dominance (FOSD) upward shift in the idiosyncratic component of the social cost parameter on the ex ante expected equilibrium actions may vary among affected players in the case of incomplete information, which is never true in the complete information case.

We next consider the effect of a denser network on equilibrium outcomes. We show that a

denser network does not necessarily lead to an increase in aggregate activity, nor does it have a monotonic effect on players' equilibrium payoffs. In particular, if we add a link between players i and j , then we show that this strictly decreases (increases) player i 's ex ante expected equilibrium action if player j 's action is less (higher) than that of player i 's social norm. More importantly, it also affects negatively (positively) a player k 's action who is path-connected to i if player j 's action is less (higher) than that of player i 's social norm.

Third, we study under which conditions policy interventions can be welfare-improving. Our analysis focuses on two main types of such interventions: those that affect players' payoffs directly, by altering the expected returns or costs associated with the activity in question, and those that affect payoffs indirectly, by changing the quality of information that is available to them or the structure of the network. Interventions of both types can be either targeted at specific players (by altering the idiosyncratic components of their payoff parameters, or by sending them a private signal) or at the general public (by altering the global components of the payoff parameters, or by sending a publicly observable signal).

Let us first focus on policies that affect players' payoffs directly. Our main question is: In order to decrease or increase to a degree ex ante expected aggregate action, should a central planner target a single player or all players in the network? In the context of crime, in attempting to maximally prevent or reduce crime, should a central planner increase the private cost to commit criminal activities for a single criminal or for all criminal in the network by, for example, increasing the probability or severity of punishment. In the context of education, should a central planner allocate scholarship funds to all students in a network or award a scholarship to a single student in order to maximally increase educational effort?

We find that, under some conditions, it may be optimal for a policy maker to concentrate their efforts in changing the payoffs of selected individuals. In the context of crime, this means that law enforcement agencies investing more resources into apprehending a key individual (*key player*)¹ may be a more efficient policy than spreading the available resources across the entire network. We also argue that unless the social planner has a good knowledge of the players' characteristics and the structure of the network, social costs are not an efficient policy instrument because of the absence of a monotonic relationship between the social cost parameters and the ex ante expected aggregate action. Such intervention may, in fact, lead to the opposite result than policy makers would like to achieve.

Fourth, we study also how the informativeness of the players' signals affects their actions and their welfare. In terms of policy, we determine to which player in the network the planner should give better information about the state of the world in order for total welfare gains from the intervention to be maximized. We show that a mean-preserving decrease (increase) in the informativeness of a player's signal about her private benefit parameter strictly decreases (increases) her ex ante expected equilibrium payoff and strictly increases (decreases) the ex ante expected equilibrium payoff of her direct neighbor.

We then consider a policy that gives more *private information* to a single player or a policy that provides more information to all players in the form of a public signal (more *public information*). We show that more private information to a single player may increase, decrease, or leave unaffected her ex ante expected equilibrium payoff and those of other offenders. A similar result may even hold true for a player whose payoff does not depend on other players' efforts. In

1. This is a different concept from the key player developed by Ballester, Calvó-Armengol, and Zenou (2006) (see Zenou (2016) for an overview) since the latter is defined as the player, who once removed, reduces aggregate activity the most.

fact, there exist values of the private benefit and social cost parameters such that more private information causes positive, negative, or no externalities on other players. The provision of a public signal may decrease ex ante expected equilibrium payoff. Therefore, more information is not always beneficial to agents.

This paper adds to the growing literature on network games,² which mostly assumes perfect information.³ There is, however, a younger literature on network games with incomplete information. Galeotti et al. (2010) develop a general network game where agents have imperfect information about the network. Calvó-Armengol and Martí (2007, 2009), Bergemann, Heumann, and Morris (2015), Martí and Zenou (2015) and Blume et al. (2015) study the linear-quadratic setting under the enrichment of a Bayesian game. Calvó-Armengol, Martí, and Prat (2015), Leister (2017) and Myatt and Wallace (2019) incorporate endogenous investment in signal precision in these settings. And in a different vein, Hagenbach and Koessler (2010) and Galeotti, Ghiglino, and Squintani (2013) study cheap-talk in networks. Finally, Golub and Morris (2017, 2018) study consistency and convergence in higher order expectations in Bayesian network games under linear best replies while Leister, Zenou, and Zhou (2019) focus on binary rather than continuous choices in a Bayesian network game.

Our paper also studies a Bayesian network game with linear best replies. Our contributions to this literature are as follows. First, in our model, the incomplete information includes all components of the players' payoff functions, in particular, private benefits and private and social costs. Second, we perform different comparative statics exercises to examine how the different parameters of the model (such as private and social costs) affects the ex ante equilibrium actions and welfare. Third, we provide different policy exercises aiming at decreasing ex ante expected aggregate action that either target a single player or all players in the network. Fourth, we study the value of information in our Bayesian network game and examine whether more information is welfare-improving and to which agent the planner should give more information. Finally, we provide an application of our model to crime under imperfect information and show how we contribute to the literature on criminal networks.

The rest of the paper is structured as follows. Section 2 illustrates all our results for criminal networks in the case of four offenders, specific network and signals. Section 3 lays out the Bayesian network game and studies its existence and uniqueness. Section 4 carries out our comparative statics analysis while the policy implications of our model are discussed in Section 5. Section 6 studies the value of information to individual players. Finally, Section 7 concludes.

The Online Appendix of this paper consists of seven parts. Section A contains the analytical derivations of the results of the criminal-network example presented in Section 2. The existence and uniqueness of the Bayesian Nash equilibrium is analyzed in Section B. Section C discusses some other numerical examples. Section D provides an alternative representation of the basic information matrix used in our analysis. A brief review of basic concepts of graph theory is provided in Section E. Some basic results of matrix analysis used in this paper have been collected in Section F. Finally, all the proofs of the statements in this paper (both in the main text and in the Online Appendix) can be found in Section G.

2. See Jackson (2008) chapter 9, Jackson and Zenou (2015) and Bramoullé and Kranton (2016) for surveys.

3. Two prominent papers in this literature are Ballester, Calvó-Armengol, and Zenou (2006) and Bramoullé, Kranton, and D'Amours (2014).

2 Criminal networks

To motivate our analysis, which is quite formal, in this section, we provide an application of our model to criminal networks for a small set of agents, for specific networks and information structure. We will illustrate not only the mechanics of our model but also all our comparative-statics and policy results. All the mathematical details of this section can be found in Section A of the Appendix.

2.1 The Bayesian network game

Following Becker (1968)'s utilitarian approach, individuals' decisions to engage in delinquent or criminal activities can be studied as an economic decision problem using a cost-benefit analysis. The decision of a potential offender depends thereby on the (expected) net payoff that she derives from an illegal activity, which is quantified as the difference between the benefits (proceeds) from the illegal activity and the costs due to the pecuniary and non-pecuniary punishment (for example, fines and incarceration) prescribed by law, weighted by the perceived probability of being apprehended.⁴ The present section revisits this topic in the context of a Bayesian network game where the potential offenders know their social or professional environment but have *incomplete information* about the cost of committing an illegal activity.

There is a finite number I of offenders. The payoff to offender $i \in \{1, \dots, I\}$ depends on the profile of nonnegative amounts of effort devoted to criminal activities, (y_1, \dots, y_I) , and an unobservable state of nature ω . It is given by:

$$u_i(\omega, (y_1, \dots, y_I)) = \bar{\alpha}_i y_i - \frac{\beta_i(\omega)}{2} y_i^2 - \frac{\bar{\gamma}_i}{2} (y_i - \bar{y}_{N_D^+(i)}})^2. \quad (1)$$

The term $\bar{\alpha}_i y_i$ represents offender i 's *private benefit* from effort y_i devoted to criminal activity. The *private benefit parameter*, $\bar{\alpha}_i$, is positive, constant, and known by all offenders. It captures offender i 's productivity or efficiency of effort in committing crime.⁵ Note that the marginal private benefit is equal to $\bar{\alpha}_i$ and does, therefore, not vary with y_i .

The term $(1/2)\beta_i(\omega)y_i^2$ represents the *private cost* inflicted on offender i . The *private cost parameter*, β_i , is a positive function of the state of nature. Its value at the state of nature ω , $\beta_i(\omega)$, is unknown to all offenders because ω is unobservable; for example, an offender may not observe the exact number of police officers on patrol in her geographical area of criminal activity, which influences the probability of apprehension. Apart from the state of nature, β_i may depend on offender i 's personal cost of engaging into criminal activity (such as time, resources, and feeling of guilt), her probability of apprehension and conviction, the severity of the statutory (pecuniary and non-pecuniary) penalties associated with the chosen effort, and the severity of the non-statutory penalties (such as shame, deterioration of future employment opportunities) resulting from a conviction. Note that the marginal private cost is equal to $\beta_i(\omega)y_i$ and is, therefore, strictly increasing in y_i .

4. In the context of network games with complete information, this has been studied by Glaeser, Sacerdote, and Scheinkman (1996), Calvó-Armengol and Zenou (2004), Ballester, Calvó-Armengol, and Zenou (2010), Liu et al. (2012), Patacchini and Zenou (2012) and Lindquist and Zenou (2014).

5. In the general Bayesian network game (see Section 3), the private benefit parameter is denoted by α_i and is unknown to all offenders. Thus, it is a function of the state of nature ω . Here, for the sake of the presentation, only the cost of illegal activity is unknown to all offenders. All variables with a bar refer to known and deterministic variables.

The term $(\tilde{\gamma}_i/2)(y_i - \bar{y}_{\mathcal{N}_D^+(i)})^2$ represents offender i 's *social cost* from deviating from the social norm of her reference group, defined as the average effort $\bar{y}_{\mathcal{N}_D^+(i)}$ of all offenders in her reference group $\mathcal{N}_D^+(i)$ in the social or professional network D . Although engaging into criminal or delinquent activities is as a rule socially condemnable, this is not true in certain social environments, such as neighborhoods with high prevalence of crime (see e.g. Kling, Ludwig, and Katz (2005) and Damm and Dustmann (2014)). As Schrag and Scotchmer (1997) put it: “gang member may quite rationally fear that his gang will punish him if he chooses not to commit a particular crime, whereas a person who is not part of a gang does not usually anticipate that his friends will penalize him if he obeys the law.” Following Akerlof (1997), the Euclidian distance between offender i 's effort and her social norm is referred to as the *social distance* between offender i and her reference group.⁶ The *social cost parameter*, $\tilde{\gamma}_i$, is nonnegative, constant, and known by all offenders; it captures offender i 's *taste for conformity*, that is, the importance she attaches to complying with her social norm, as well as the strength of that norm in her social environment.⁷ Compared to offenders with lower social cost parameters, offenders with larger social cost parameters incur higher social costs from deviating from their social norms because they are stronger in their neighbourhood or society or because adhering to the social norm of their reference group is intrinsically more important to these offenders. Note that the marginal social cost is equal to $\tilde{\gamma}_i(y_i - \bar{y}_{\mathcal{N}_D^+(i)})$ and is, therefore, strictly increasing in y_i and strictly decreasing in $\bar{y}_{\mathcal{N}_D^+(i)}$ if $\tilde{\gamma}_i > 0$.

In order to keep the exposition simple, we consider four offenders ($I = 4$) whose social or professional network is represented by the digraph D (directed network) depicted in Figure 2.1; therein offenders are represented by vertices and social or professional ties by arrows, called arcs or directed links. An directed link from offender i to offender j is denoted by (i, j) and signifies that offender j belongs to the reference group $\mathcal{N}_D^+(i)$ of offender i , in which case offender j 's effort enters offender i 's social norm $\bar{y}_{\mathcal{N}_D^+(i)}$ and in turn affects her payoff, which in turn determines her own effort. A directed link (i, j) (outdegree) can be interpreted as offender i paying attention to the behavior of offender j or, alternatively, offender j being a role model to offender i . It follows from Figure 2.1 that offender 1's reference group consists of offender 2, offender 2's reference group of offender 3, offender 3's reference group of offender 2, and offender 4's reference group of offenders 1 and 3, that is

$$\mathcal{N}_D^+(1) = \{2\}, \quad \mathcal{N}_D^+(2) = \{3\}, \quad \mathcal{N}_D^+(3) = \{2\}, \quad \mathcal{N}_D^+(4) = \{1, 3\}.$$

The social norms are therefore given by

$$\bar{y}_{\mathcal{N}_D^+(1)} = y_2, \quad \bar{y}_{\mathcal{N}_D^+(2)} = y_3, \quad \bar{y}_{\mathcal{N}_D^+(3)} = y_2, \quad \bar{y}_{\mathcal{N}_D^+(4)} = \frac{y_1 + y_3}{2}.$$

The four potential offenders reside in one of two precincts of a city: precinct East and precinct West; specifically, offenders 1 and 2 reside in precinct East and offenders 3 and 4 in precinct West. All offenders are criminally active only in their precinct of residence. A city authority is

6. This is the standard way economists have modeled conformism in networks. See, in particular, Patacchini and Zenou (2012), Boucher (2016), Sommer (2017) and Ushchev and Zenou (2019), for the case of perfect information and Calvó-Armengol and Martí (2009), Calvó-Armengol, Martí, and Prat (2015) and Golub and Morris (2018) for the case of imperfect information.

7. In the general Bayesian network game (see Section 3), the social cost parameter is denoted by γ_i and is a function of the state of nature ω .

responsible for the allocation of financial and human resources to policing functions, including the geographical allocation of police forces between precincts East and West.

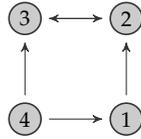


Figure 2.1 Network D of offenders

A key element affecting an individual's decision about whether, and to what extent, to commit a crime is her *perception of the probability of being apprehended and convicted*. There exists a rich literature on how these perceptions are formed and on how accurate they are. A multitude of factors may affect these perceptions, among others, own age (Gardner and Steinberg 2005; Hjalmarsson 2008), own experience (Lochner 2007), experience of peers (Parker and Grasmick 1979), and the prevalence of similar crimes in the area or jurisdiction (Schrag and Scotchmer 1997).⁸ In our model, an offender's private cost is strictly increasing in the probability of apprehension and conviction. For a particular offender, this probability depends on the fraction of police officers on duty in her precinct of residence (at the time a crime is committed) and the strictness of the judge or jury deciding her case if she is apprehended and prosecuted. Thus, the probability of apprehension and conviction is increasing in these two factors. Indeed, the more policemen are patrolling the streets or working on solving criminal cases in her precinct of residence, the more likely an offender is to be brought to justice and convicted. Also, the stricter is the judge or the jury hearing her case, the more likely is a conviction and the harsher is the expected punishment.

The fraction of police officers on duty in precinct East, denoted by $\omega_1 \in [0, 1]$, and a measure of the strictness of the judge or jury, denoted by $\omega_2 \in [0, 1]$, constitute the state of nature ω , that is, $\omega = (\omega_1, \omega_2) \in [0, 1]^2$. The relation between these two dimensions of the state of nature and offender i 's private cost parameter depends on her precinct of residence. Two possible cases are depicted in Figure 2.2 and based on the following assumption: an offender's private cost parameter is either low ($\theta_{\beta, \text{low}}$) or high ($\theta_{\beta, \text{high}}$); it is high if and only if both the fraction of police officers on duty in her precinct of residence and the measure of the strictness of the judge or jury are not less than one half.

If she resides in precinct East, then $\beta_i((\omega_1, \omega_2)) = \theta_{\beta, \text{high}}$ if and only if $\omega_1 \geq 1/2$ and $\omega_2 \geq 1/2$ (see the right panel of Figure 2.2; therein, the area shaded in dark gray is the set of states of nature on which the private cost parameter assumes the high value). If the offender resides in precinct West, $\beta_i((\omega_1, \omega_2)) = \theta_{\beta, \text{high}}$ if and only if $\omega_1 \leq 1/2$ and $\omega_2 \geq 1/2$ (see the left panel of Figure 2.2).

Prospective offenders cannot observe the state of nature, which determines the values of their private cost parameters. Indeed, they don't know the exact fraction of police officers on duty in the precinct where they live and they don't know in advance which judge will hear their case or the composition of the jury if they are apprehended and prosecuted. They have only incomplete information about these two dimensions of the state of nature, based on their observations, their experience, and the experiences of others. The incomplete information of offender i about the state of nature is formally represented by a function $s_{\beta, i}$, called signal (component), whose value only depends on the state of nature ω .⁹ There are only two possible values that the signal $s_{\beta, i}$ can

8. For literature reviews, see Apel (2013) and Chalfin and McCrary (2017).

9. In the general Bayesian network game (see Section 3), where the private benefit, the private cost, and the social

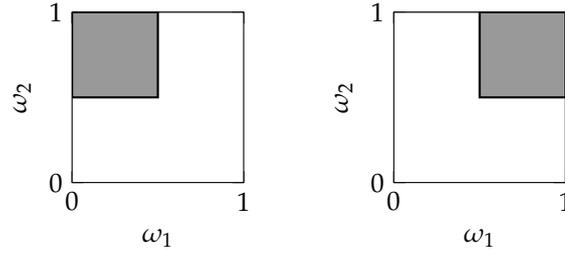


Figure 2.2 Private cost parameters of offenders residing in precinct West (left panel) and precinct East (right panel). Areas shaded in dark gray represent states of nature that give rise to a high value of the private cost parameter ($\theta_{\beta,\text{high}}$).

take, low or high, which represent private information, that is, they are not observable by other offenders. An example is depicted in Figure 2.3; therein, the area shaded in dark gray is the set of states of nature on which the signal assumes the high value. By observing a high signal value, the offender learns that the state of nature lies in the dark shaded area, and by observing a low signal value, she learns that it lies in the white area. Although the signal value does not reveal the state of nature, it conveys information about the value of the private cost parameter at the state of nature.

Let us consider two cases. First, suppose offender i resides in precinct East. Her signal $s_{\beta,i}$ (given in Figure 2.3) is completely informative about β_i because it reveals the value of β_i (see the right panel of Figure 2.2). Loosely speaking, the signal is correct about the two dimensions of the state of nature. Second, suppose offender i resides in precinct West. Her signal $s_{\beta,i}$ is only partially informative about β_i (see the left panel of Figure 2.2). Indeed, if the value of $s_{\beta,i}$ is high, she learns that the value of β_i is in fact $\theta_{\beta,\text{low}}$. If the value of $s_{\beta,i}$ is low, she learns that the value of β_i is either $\theta_{\beta,\text{low}}$ or $\theta_{\beta,\text{high}}$. Loosely speaking, the signal is misleading about the first dimension (police) and correct about the second dimension (judge or jury) of the state of nature.

The preceding example provides an important insight: the informational content of a signal is embodied in the partition it induces on the set of all possible states of nature (which is the unit square $[0, 1]^2$). The actual values a signal can take, referred to as low and high, are immaterial, provided that they are different. It is therefore natural to identify the low value with $\theta_{\beta,\text{low}}$ and the high value with $\theta_{\beta,\text{high}}$.

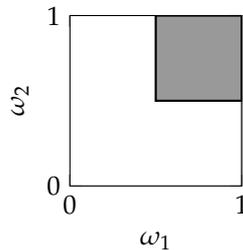


Figure 2.3 Example of private signal. Area shaded in dark gray represents states of nature that give rise to a high value of the signal ($\theta_{\beta,\text{high}}$).

cost parameters are *all* functions of the state of nature, a player's signal consists of three components; more specifically, player i ' signal is denoted by s_i and consists of the components $s_{\alpha,i}, s_{\beta,i}, s_{\gamma,i}$, that is, $s_i := (s_{i,\alpha}, s_{i,\beta}, s_{i,\gamma})$.

Let us go back to our initial example with four offenders connected by the network given in Figure 2.1. The private cost parameters and the signals of the four offenders are depicted at the top and bottom panels of Figure 2.4. The first column corresponds to offender 1, the second column to offender 2, etc. Recall that offenders 1 and 2 reside in precinct East and offenders 3 and 4 in precinct West. The values of their private cost parameters depend on their precinct of residence as described in Figure 2.2.¹⁰ Assuming that the state of nature follows a uniform distribution, then, their prior beliefs are such that the probability of having a high-cost parameter is $1/4$, i.e., $\mathbb{P}(\theta_{\beta,\text{high}}) = 1/4$ while that of having a low-cost parameter parameter is $3/4$, i.e., $\mathbb{P}(\theta_{\beta,\text{low}}) = 3/4$. Let us now calculate their posterior beliefs after receiving their signals.

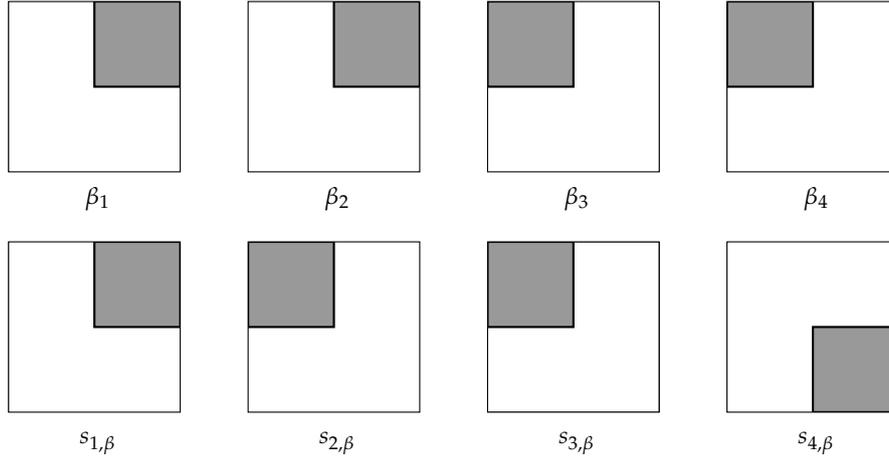


Figure 2.4 Offenders' private cost parameters (top panels) and private signals (bottom panels)

Contrasting offender 1's signal $s_{1,\beta}$ with her private cost parameter β_1 shows that $s_{1,\beta}$ is completely informative about β_1 . There is no uncertainty about the value of β_1 because $s_{1,\beta}$ reveals its value for all possible states of nature.¹¹ Loosely speaking, the signal is correct about the two dimensions of the state of nature. Contrasting offender 2's signal $s_{2,\beta}$ with her private cost parameter β_2 shows that $s_{2,\beta}$ is partially informative about β_2 . There is no uncertainty about the value of β_2 at states of nature that give rise to a high signal value (the area shaded in dark gray) since offender 2 knows with certainty that the state is low,¹² and there is uncertainty about the value of β_2 at states of nature that give rise to a low signal value (the white area). Using Bayes' law and assuming that the state of nature follows a uniform distribution, by observing the value $\theta_{\beta,\text{low}}$ for $s_{2,\beta}$ offender 2 learns that β_2 is low ($\theta_{\beta,\text{low}}$) with probability $2/3$ and high

10. In Appendix A.1, we formally derive the private cost parameters and the signals depicted in Figure 2.4.

11. Indeed, by using Bayes' law and assuming that the state of nature follows a uniform distribution, it is easily verified that: $\mathbb{P}(\theta_{\beta,\text{low}}|s_{\beta,\text{low}}) = \mathbb{P}(\theta_{\beta,\text{high}}|s_{\beta,\text{high}}) = 1$ and $\mathbb{P}(\theta_{\beta,\text{high}}|s_{\beta,\text{low}}) = \mathbb{P}(\theta_{\beta,\text{low}}|s_{\beta,\text{high}}) = 0$.

12. Indeed, by using Bayes' law and assuming that the state of nature follows a uniform distribution, we have:

$$\begin{aligned} \mathbb{P}(\theta_{\beta,\text{low}}|s_{\beta,\text{high}}) &= \frac{\mathbb{P}(s_{\beta,\text{high}}|\theta_{\beta,\text{low}})\mathbb{P}(\theta_{\beta,\text{low}})}{\mathbb{P}(s_{\beta,\text{high}}|\theta_{\beta,\text{low}})\mathbb{P}(\theta_{\beta,\text{low}}) + \mathbb{P}(s_{\beta,\text{high}}|\theta_{\beta,\text{high}})\mathbb{P}(\theta_{\beta,\text{high}})} \\ &= \frac{1/3 \times 3/4}{1/3 \times 3/4 + 0 \times 1/4} = 1 \end{aligned}$$

And, $\mathbb{P}(\theta_{\beta,\text{high}}|s_{\beta,\text{high}}) = 1 - \mathbb{P}(\theta_{\beta,\text{low}}|s_{\beta,\text{high}}) = 0$.

$(\theta_{\beta,\text{high}})$ with probability $1/3$.¹³ As a result, starting with the prior that the probability of having a low cost (high cost) of committing crime, i.e., $\theta_{\beta,\text{low}}$ ($\theta_{\beta,\text{high}}$) is $3/4$ ($1/4$), offender 2 updates her beliefs as follows: if she receives a low signal, the probability of $\theta_{\beta,\text{low}}$ is $2/3$ and that of $\theta_{\beta,\text{high}}$ is $1/3$. If she receives a high signal, then she knows with certainty that the state is low.

Contrasting offender 3's signal $s_{3,\beta}$ with her private cost parameter β_3 shows that $s_{3,\beta}$ is completely informative about β_3 . Finally, contrasting offender 4's signal $s_{4,\beta}$ with her private cost parameter β_4 shows that $s_{4,\beta}$ is partially informative about β_4 . Loosely speaking, the signal is misleading about both dimensions of the state of nature. There is, however, no uncertainty about the value of β_4 at states of nature that give rise to a high signal value, and there is uncertainty about the value of β_2 at states of nature that give rise to a low signal value. By observing the value $\theta_{\beta,\text{low}}$ for $s_{4,\beta}$ offender 4 learns that β_4 is low ($\theta_{\beta,\text{low}}$) with probability $2/3$ and high ($\theta_{\beta,\text{high}}$) with probability $1/3$. In summary, the four pairs of private cost parameters and signals differ in some respects; they have in common that a high signal value reveals the value of the private cost parameter, that is, an offender learns the value of her private cost parameter by observing a high signal value.

2.2 The Bayesian Nash equilibrium

Let us now determine the Bayesian Nash equilibrium of this game. Each offender draws up a rule, called strategy, specifying the amount of effort to devote to criminal activity depending on the value of her signal. In our game of crime, where the network connecting offenders is D (Figure 2.1), offender i 's strategy is denoted by $x_i(D): \{\theta_{\beta,\text{low}}, \theta_{\beta,\text{high}}\} \rightarrow \mathbb{R}_+$. A profile of strategies $(x_1^*(D), \dots, x_4^*(D))$ is a *Bayesian Nash equilibrium* in the game of crime if each of its strategies benefits each agent the most (in terms of ex ante expected payoff) and no offender has an incentive to unilaterally deviate from her strategy, in which case the strategies in the aforementioned profile are called *equilibrium strategies*. Offender i 's equilibrium strategy $x_i^*(D)$ at the value $\vartheta_i \in \{\theta_{\beta,\text{low}}, \theta_{\beta,\text{high}}\}$ of her signal $s_{i,\beta}$, $x_i^*(D)(\vartheta_i)$, is characterized by the equality of her marginal private benefit, $\bar{\alpha}_i$, and her predicted marginal cost, defined as the sum of predicted private marginal cost, $x_i^*(D)(\vartheta_i) \mathbb{E}(\beta_i | s_{i,\beta} = \vartheta_i)$, and predicted social cost, $\bar{\gamma}_i(x_i^*(D)(\vartheta_i) - \mathbb{E}(\bar{x}_{\mathcal{N}_D^+(i)}^*(D) | s_{i,\beta} = \vartheta_i))$, where $\bar{x}_{\mathcal{N}_D^+(i)}^*(D)$ denotes her social norm at the equilibrium. So, clearly, in equilibrium, each offender i has to form expectations on her marginal cost of committing crime, i.e., $\mathbb{E}(\beta_i | s_{i,\beta} = \vartheta_i)$ and on the cost of committing crime of her direct links because it affects her social norm, i.e., $\mathbb{E}(\bar{x}_{\mathcal{N}_D^+(i)}^*(D) | s_{i,\beta} = \vartheta_i)$.¹⁴

Although a formal presentation of this game is deferred to Section 3, we can see how offenders interact and affect each other within the context of our example.¹⁵ Suppose that the state of

13. Indeed,

$$\begin{aligned} \mathbb{P}(\theta_{\beta,\text{low}} | s_{\beta,\text{low}}) &= \frac{\mathbb{P}(s_{\beta,\text{low}} | \theta_{\beta,\text{low}}) \mathbb{P}(\theta_{\beta,\text{low}})}{\mathbb{P}(s_{\beta,\text{low}} | \theta_{\beta,\text{low}}) \mathbb{P}(\theta_{\beta,\text{low}}) + \mathbb{P}(s_{\beta,\text{low}} | \theta_{\beta,\text{high}}) \mathbb{P}(\theta_{\beta,\text{high}})} \\ &= \frac{2/3 \times 3/4}{2/3 \times 3/4 + 1 \times 1/4} = \frac{2}{3} \end{aligned}$$

And, $\mathbb{P}(\theta_{\beta,\text{high}} | s_{\beta,\text{low}}) = 1 - \mathbb{P}(\theta_{\beta,\text{low}} | s_{\beta,\text{low}}) = \frac{1}{3}$.

14. Observe that the signals of the different players are correlated through the state variable. For example, assuming that the state of nature follows a uniform distribution, it is easily verified that: $\mathbb{P}(s_{4,\beta} = \theta_{\beta,\text{low}} | s_{3,\beta} = \theta_{\beta,\text{low}}) = 2/3$, $\mathbb{P}(s_{4,\beta} = \theta_{\beta,\text{high}} | s_{3,\beta} = \theta_{\beta,\text{low}}) = 1/3$, $\mathbb{P}(s_{4,\beta} = \theta_{\beta,\text{low}} | s_{3,\beta} = \theta_{\beta,\text{high}}) = 1$ and $\mathbb{P}(s_{4,\beta} = \theta_{\beta,\text{high}} | s_{3,\beta} = \theta_{\beta,\text{high}}) = 0$.

15. Proposition B.1 establishes existence and uniqueness of a Bayesian Nash equilibrium for the general case, and provides its functional form. For a more detailed derivation of the numerical values of the present example, see

nature follows a uniform distribution, $\bar{\alpha}_1 = \bar{\alpha}_2 = \bar{\alpha}_3 = \bar{\alpha}_4 = 1$, $\theta_{\beta,low} = 1/2$ and $\theta_{\beta,high} = 3/2$, and $\bar{\gamma}_1 = \bar{\gamma}_2 = 1/2$ and $\bar{\gamma}_3 = \bar{\gamma}_4 = 1/3$. Then, at the Bayesian-Nash equilibrium, the offenders' ex ante expected efforts and payoffs in equilibrium, as well as the social norms each one expects to face, are presented in Table 2.1.¹⁶

Table 2.1 Ex ante expected equilibrium efforts, payoffs, and social norms

offender	effort	payoff	social norm
1	1.49	0.77	1.40
2	1.40	0.62	1.52
3	1.52	0.79	1.40
4	1.42	0.76	1.51

Notes: Real numbers in decimal notation are rounded to two decimal places.

We can see that, in equilibrium, offender 3 exerts the highest effort, has the highest payoff but faces the lowest social norm so that her effort is above her social norm. On the contrary, offender 2 makes the lowest effort, obtains the lowest payoff and faces the highest social norm so that her effort is below her social norm. This is interesting because both offenders 2 and 3 are ex ante identical and have exactly the same position in the network and their social norm only consists of each other. The only differences between the two of them are their location (offender 2 resides in precinct East while offender 3 lives in precinct West) and the signals they receive. As stated above, offender 2 receives a noisy signal. In particular, when she receives a high signal she knows with certainty that the state of the world is low, i.e., $\mathbb{E}(\beta_2 | s_{2,\beta} = \theta_{\beta,high}) = \theta_{\beta,low}$, while, when she receives a low signal, she knows with probability 2/3 that the cost of committing crime is low (i.e. less police and less severe judges) and with probability 1/3 that the cost of committing crime is high, i.e., $\mathbb{E}(\beta_2 | s_{2,\beta} = \theta_{\beta,low}) = \frac{2}{3}\theta_{\beta,low} + \frac{1}{3}\theta_{\beta,high}$. On the contrary, offender 3 receives a perfect informative signal so that a low signal indicates with probability 1 that the state is low and the same for a high signal. This is the expectation each offender has on herself. However, they also need to form expectations on the criminal effort of the other offenders because it affects their social norm. We can understand why offender 2 (but also offender 4) exerts a low effort. It is because she receives a noisy signal; with probability 1/3 she wrongly believes that the cost of committing crime is high when she receives a low signal. Since offender 3 knows this, she expects a low social norm. The reverse is true for offender 2.

This model just shows that imperfect information about the presence of police in one's neighborhood and the severity of the judge if someone is arrested strongly affect offenders' behavior and should be taken into account if one wants to address policies aiming at reducing crime. We will now consider three types of policy interventions: policies aiming at changing the payoff parameters, policies aiming at changing the network structure, and policies aiming at changing the availability or quality of information.

2.3 Changing payoff parameters

Perhaps the most intuitive policy aimed towards reducing overall criminal activity is through interventions that increase offenders' cost of engaging into crime (parameter β). In the context

Section A in the Appendix.

16. For a formal derivation, see Section A.1 of the Appendix.

of our model, examples of such policies are an increase in the size of the police force and/or the introduction of more severe statutory penalties for offenders. The effect of changes in the offenders' payoffs parameters in the general case is discussed in Section 4.1, and the relevant policy implications in Section 5.2.

To discuss the effects of changes in the private cost parameters, consider a city authority that decides to increase their budget for law enforcement, with the objective of reducing criminal activity at the city level by increasing the likelihoods of arresting perpetrators. This raises the question about which allocation of the additional resources is most effective in reducing criminal activity. Two policy alternatives are considered: targeting a single offender, in which case all additional resources are devoted to her apprehension, or targeting all offenders, in which case the resources are spread across the entire network in order to, for example, increase uniformly the probability of apprehension. As our discussion of optimal targeting (see Section 5.2) shows, neither alternative is superior to the other. We now analyze these two policies in the context of our specific example.

Suppose the additional resources suffice either to increase the private cost parameter of a single offender by some positive amount $\Delta\bar{\beta}$ (first policy) or to increase uniformly the private cost parameters of all four offenders by $(1/4)\Delta\bar{\beta}$ (second policy). Table 2.2 reports changes in ex ante expected aggregate equilibrium effort (criminal activity) for the aforementioned two policies for two different networks, D_1 and D_4 (network D_4 is clearly denser and more connected than D_1)¹⁷ and two different values of $\Delta\bar{\beta}$, $2/3 = 0.666$ and $5/6 = 0.833$. In the game of crime with network D_4 , targeting all offenders is superior to targeting a single offender because it reduces ex ante expected aggregate criminal activity most, irrespective of the value of $\Delta\bar{\beta}$. By contrast, in the game of crime with network D_1 , which of the two policies is more effective depends on the value of $\Delta\bar{\beta}$: in case $\Delta\bar{\beta} = 2/3$, targeting a single player is superior to targeting all players; in case $\Delta\bar{\beta} = 5/6$, targeting all offenders is superior to targeting a single offender. This shows that the effectiveness of each policy alternative depends on the network structure as well as the size of the available resources.

Table 2.2 Changes in ex ante expected aggregate effort

Targeted offender(s)	$\Delta\bar{\beta} = 2/3$		$\Delta\bar{\beta} = 5/6$	
	D_1	D_4	D_1	D_4
1	-0.655	-0.855	-0.747	-0.969
2	-1.029	-0.988	-1.173	-1.126
3	-1.213*	-0.958	-1.358	-1.078
4	-0.556	-0.694	-0.632	-0.788
{1, 2, 3, 4}	-1.182	-1.185*	-1.406*	-1.409*

Notes: Real numbers in decimal notation are rounded to three decimal places. Numbers with a star indicate a column minimum.

Given the above finding, an alternative policy aimed at changing the payoff parameters could focus on changing the value of $\tilde{\gamma}_i$, the taste for conformity, capturing the offender's social cost from deviating from the norm of her reference group (see (1)). In particular, in crime-ridden communities, reducing the social norm of crime could be effective in reducing total crime.

17. See Figure 2.5 for a graphical representation where network D_1 corresponds to the network described in Figure 2.1.

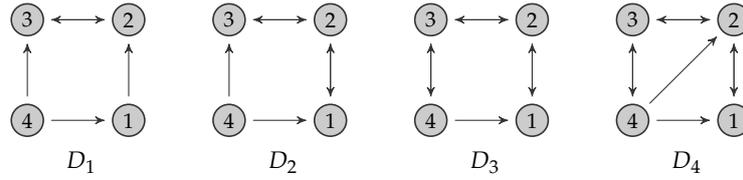


Figure 2.5 Evolution of offenders' network over time

Our general analysis (see Result 6 of Proposition 4.6) shows that a uniform decrease in the offenders' strengths of their social norms, as represented by their social cost parameters, has, in general, an ambiguous effect on reducing crime (see also Tables A.1 and A.2). An alternative to the above policy measure consists in targeting only a single offender or a group of offenders, for whom conforming to a social norm is intrinsically important. Our analysis (see Result 5 of Proposition 4.6 and the example in Section A.2 in the Appendix) shows that lowering the strength of the social norm of a single offender may decrease, increase, or leave unchanged her ex ante expected equilibrium effort. The sign of the effect depends on the structure of the network and the offenders' signals and payoff parameters. A policy intervention aimed at changing the social norms of a group of offenders must, therefore, be tailored to the network and the offenders' personal characteristics. An individual-tailored policy could, for example, target younger gang members who—in spite of having shown high ability in areas such as education, sports, or arts—are dragged into criminal activities in order to avoid social exclusion, ostracization, or even violence. Offering college scholarships to such individuals could provide them with better future prospects, reduce their benefits from conforming to the norm of their reference group, and make their costs from engaging into delinquent behavior more salient to them. A notable example is the case of the Comer school, mentioned by Akerlof (1997).¹⁸

2.4 Changing the network

A set of policies that can be used to combat crime are those that focus on disrupting the network of offenders by cutting their communication and their influence channels. Such interventions are formally studied in Sections 4.2 and 5.1. A rather interesting and perhaps counter-intuitive finding is that severing directed links, even between relatively active offenders, does not necessarily lead to a decrease in the aggregate level of crime. This implies that, under relatively weak assumptions, a dense criminal network may be associated with a lower overall crime rate than a sparser network (see Proposition 4.8).

To discuss the effects of changes in the topology of the network connecting offenders, we consider a network formation game over four periods of time. In the first period, the network is D_1 . In every period $\tau \in \{2, \dots, 4\}$, offender τ expands her reference group by forming a single directed link to another offender such that the new directed link benefits her most (in terms of ex ante expected equilibrium payoff). The resulting sequence of networks D_1, D_2, D_3, D_4 is displayed in Figure 2.5. In network D_4 , each offender's payoff is affected by the choices of efforts of other offenders and each offender's choice of effort affects other offenders' payoffs, either directly or indirectly. By contrast, in network network D_1 , offender 4's choice of effort does not

18. The effects of changes in one of the three kinds of payoff parameters (the private benefit parameters, the private cost parameters, or the social cost parameters) in the general case are discussed in Section 4.1 and the relevant policy implications in Section 5.2.

affect any of the other offenders' payoffs.

Figure 2.6 displays ex ante expected equilibrium efforts (left panels), social distances (middle panels), and payoffs (right panels) of offenders 2 (top panels) and 3 (bottom panels) in period τ where the network connecting offenders is D_τ . The comparison of values for two consecutive periods shows the effects of a new directed link on the above three quantities.¹⁹ The following discussion focuses on periods 1 and 2 and offenders 2 and 3 (remember that offender 3 receives signals that are perfectly informative while offender 2 receives noisy signals).

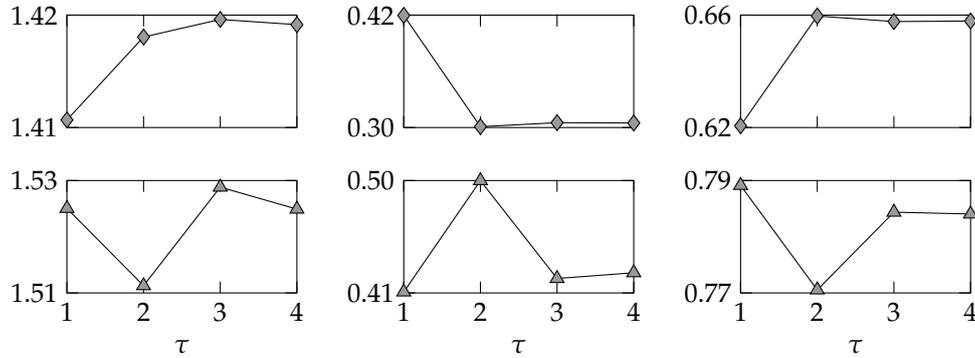


Figure 2.6 Ex ante expected equilibrium efforts (left panels), social distances (middle panels), and payoffs (right panels) of offenders 2 (\diamond , upper panel) and 3 (\blacktriangle , lower panel) in the game of crime in period τ

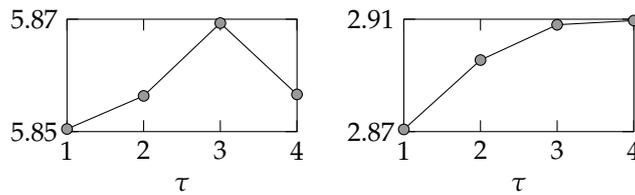


Figure 2.7 Ex ante expected aggregate equilibrium effort (left panel) and welfare (right panel) in the game of crime in period τ

In period 2, offender 2 forms a new directed link to offender 1. As will be explained below, this increases her ex ante expected equilibrium effort and decreases that of her role model, offender 3.²⁰ The signs of the effects vary, therefore, among the two offenders. We show (see Proposition 4.9) that a variation in the signs of such effects is not possible in the case of complete information, where there is no uncertainty about payoffs. This result suggests that an increase in uncertainty dampens the impact of exogenous shocks to the topology of the network at the macro level in terms of a smaller variation in aggregate (criminal) activity. Note that the magnitude of the effect is larger for offender 3 than offender 2.²¹

This shows that there is in general no effect attenuation within the network; specifically, the magnitude of the change in ex ante expected equilibrium effort does not decrease with the distance to the offender who is the source of the change (offender 2).

19. See Section A.3 in the Appendix for results on offenders 1 and 4.

20. All numerical values relevant to the present discussion can be found in Section A.3 in the Appendix.

21. Offender 2's ex ante expected equilibrium effort increases by 0.0089 and that of offender 3 decreases by 0.0096 (numbers rounded to four decimal places).

By design of the network formation game, offender 2 chooses her new role model among all offenders in order to maximally increase her ex ante expected equilibrium effort. This explains the increase in offender 2's ex ante expected equilibrium payoff between periods 1 and 2 (see the top right panel in Figure 2.6). She achieves this increase by reducing the ex ante expected social distance to her reference group (see the top middle panel in Figure 2.6), defined as the Euclidian distance between her effort and her social norm, for an offender's payoff is strictly decreasing in her social distance to her reference group. Offender 2 is the only role model of offender 3. The changes in offender 2's equilibrium behavior (that is, the changes in her equilibrium efforts at the low and high value of her signal, which cause an increase in her ex ante expected equilibrium effort), represent, therefore, changes in the social norm of offender 3. These changes induce offender 3 to change her equilibrium behavior, which entail a decrease in her ex ante expected equilibrium effort and payoff. This decrease in her ex ante expected equilibrium payoff represent an externality of offender 2's formation of a new directed link to offender 1. Note that such externalities may be negative (as in the case of offender 3) or positive (as in the cases of offenders 1 and 4).²²

In summary, there is no monotone relation between the density of the network connecting offenders and ex ante expected equilibrium effort, social distance, or payoff. The left panel of Figure 2.7 shows that ex ante expected aggregate equilibrium effort, that is, aggregate criminal activity, is also not monotone in the network density. A policy measure to decrease network density (for example, the isolation of a convicted offender from her reference group by enforcing association, location, or residence restrictions) is, therefore, in general not always expedient to reduce aggregate criminal activity. Finally, note that ex ante expected equilibrium welfare is strictly increasing in the network density, as shown in the right panel of Figure 2.7.²³

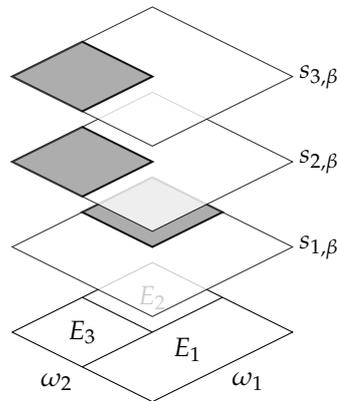


Figure 2.8 Partition of set of states of nature $[0, 1]^2$ generated by signals $s_{1,\beta}, s_{2,\beta}, s_{3,\beta}: \{E_1, E_2, E_3\}$

The concept of equilibrium strategy described above is the key to understanding the mechanism underlying these results. For example, let us explain why offender 2's ex ante expected equilibrium effort increases in period 2 when she forms a new directed link to offender 1. The reference group of offender 2 in period 1 consists of offender 3 and in period 2 of offenders 1 and 3. We need, therefore, to determine all possible combinations of values of the signals of offenders 1, 2, and 3 that may occur together. According to the description of the game of

22. See Section A.3 in the Appendix, Figure A.5 in particular.

23. This monotonicity is rooted in the design and the structure of the network formation game. In general, this is not true.

crime, if $s_{1,\beta} = \theta_{\beta,\text{low}}$, then either $s_{2,\beta} = s_{3,\beta} = \theta_{\beta,\text{low}}$ or $s_{2,\beta} = s_{3,\beta} = \theta_{\beta,\text{high}}$, and if $s_{1,\beta} = \theta_{\beta,\text{high}}$, then $s_{2,\beta} = s_{3,\beta} = \theta_{\beta,\text{low}}$. These three combinations generate a partition of the set of states of nature $[0, 1]^2$ into three events,

$$\begin{aligned} E_1 &:= \{s_{1,\beta} = \theta_{\beta,\text{low}}, s_{2,\beta} = \theta_{\beta,\text{low}}, s_{3,\beta} = \theta_{\beta,\text{low}}\}, \\ E_2 &:= \{s_{1,\beta} = \theta_{\beta,\text{high}}, s_{2,\beta} = \theta_{\beta,\text{low}}, s_{3,\beta} = \theta_{\beta,\text{low}}\}, \\ E_3 &:= \{s_{1,\beta} = \theta_{\beta,\text{low}}, s_{2,\beta} = \theta_{\beta,\text{high}}, s_{3,\beta} = \theta_{\beta,\text{high}}\}, \end{aligned}$$

which are depicted in Figure 2.8. The values of the three offenders' equilibrium strategies that are relevant for the present discussion are reported in Table 2.3.²⁴

Table 2.3 Selected values of offender 1, 2, and 3's equilibrium strategies in the games of crime in periods 1 and 2

	Value of $(\vartheta_1, \vartheta_2, \vartheta_3)$		
	$(\theta_{\beta,\text{low}}, \theta_{\beta,\text{low}}, \theta_{\beta,\text{low}})$	$(\theta_{\beta,\text{high}}, \theta_{\beta,\text{low}}, \theta_{\beta,\text{low}})$	$(\theta_{\beta,\text{low}}, \theta_{\beta,\text{high}}, \theta_{\beta,\text{high}})$
$x_1^*(D_1)(\vartheta_1)$	1.704	0.853	1.704
$x_2^*(D_1)(\vartheta_2)$	1.412	1.412	1.4
$x_3^*(D_1)(\vartheta_3)$	1.765	1.765	0.8
$x_2^*(D_2)(\vartheta_2)$	1.343	1.343	1.641

Notes: Real numbers in decimal notation with more than one digit after the decimal mark are rounded to three decimal places.

First, we discuss changes of quantities taking place on the event for which offender 2's signal assumes a low value, that is, the event $\{s_{2,\beta} = \theta_{\beta,\text{low}}\} = E_1 \cup E_2$ (see the second and third columns of Table 2.3). In period 1, offender 2's effort (1.412) is less than her social norm, offender 3's effort (1.765), so that her marginal social cost is negative ($-0.176 = (1/2)(1.412 - 1.765)$). Evidently, her marginal private benefit (1) is equal to the sum of her predicted marginal private cost ($1.176 = 1.412((1/2)(2/3) + (3/2)(1/3) = 1.412 \mathbb{E}(\beta_2 | s_{2,\beta} = \theta_{\beta,\text{low}}))$) and her marginal social cost (-0.176). By forming a directed link to offender 1 in period 2, offender 2's predicted social norm decreases (from 1.765 to $1.593 = (1.765 + 1.420)/2$) because her prediction of offender 1's effort ($1.420 = 1.704(2/3) + 0.853(1/3)$) is less than that of offender 3's effort (1.765). As a result, offender 2's predicted marginal social cost increases (from -0.176 to $-0.090 = (1/2)(1.412 - 1.593)$), so that her marginal private benefit (1) is less than the sum of her predicted marginal private cost (1.176) and her predicted marginal social cost (-0.090). In order to correct this imbalance, offender 2 must decrease her effort on the event $\{s_{2,\beta} = \theta_{\beta,\text{low}}\}$ (because predicted marginal private cost is strictly increasing in effort); in fact, her effort at the low signal value decreases from 1.412 to 1.343.

Second, we discuss changes of quantities taking place on the event $\{s_{2,\beta} = \theta_{\beta,\text{high}}\} = E_3$ (see the last column of Table 2.3). In period 1, offender 2's effort (1.4) is larger than her social norm, offender 3's effort (0.8), so that her marginal social cost is positive ($0.3 = (1/2)(1.4 - 0.8)$). Evidently, her marginal private benefit (1) is equal to the sum of her predicted marginal private cost ($0.7 = 1.4(1/2) = 1.4 \mathbb{E}(\beta_2 | s_{2,\beta} = \theta_{\beta,\text{high}})$) and her marginal social cost (0.3). The new directed link to offender 1 increases her social norm (from 0.8 to $1.252 = (1.704 + 0.8)/2$). As a result, offender 2's marginal social cost decreases (from 0.3 to $0.074 = (1/2)(1.4 - 1.252)$), so

24. The values are taken from Section A.3 in the Appendix.

that her marginal private benefit (1) is larger than the sum of her predicted marginal private cost (0.7) and marginal social cost (0.074). In order to correct this imbalance, offender 2 must increase her effort on the event $\{s_{2,\beta} = \theta_{\beta,\text{high}}\}$; in fact, her effort at the high signal value increases from 1.4 to 1.641.

It follows from the above that offender 2's ex ante expected equilibrium effort increases from $1.409 = 1.412(3/4) + 1.4(1/4)$ to $1.418 = 1.343(3/4) + 1.641(1/4)$.

2.5 Changing the informativeness of signals

An alternative approach would be to promote policies aimed at changing the perceptions that prospective offenders have about their cost from engaging into criminal activity. As discussed in Section 2.1, offenders are not fully cognizant of their probability of apprehension and conviction, and can only make inferences based on the signal values they observe. A question that naturally arises in that case is how a *ceteris paribus* more informative signal would affect aggregate criminal activity, and whether criminals would be better or worse off with more information. This is discussed in detail for the general model in Section 6.

Baumann and Friehe (2013) develop a model in which the policy maker may choose to costlessly communicate information about the probability of conviction to the potential offenders. They find that better information is in general beneficial for offenders, even though it may also lead to lower the crime rates for certain parameter values. In the present model, however, in the presence of peer effects and positive social costs, this is not as clear. A more informative signal can push aggregate criminal activity in either direction and the same holds true for the welfare of the offenders. The question we ask here is whether better information about the state of the world may be welfare-improving in network models where social norms matter.

Let us now study this policy in a simple example. For the purposes of obtaining nontrivial results, its structure will be slightly modified.²⁵ Specifically, an offender's private cost parameter can *now* have three different values: $\theta_{\beta,\text{low}}$, $\theta_{\beta,\text{mid}} := (1/2)(\theta_{\beta,\text{low}} + \theta_{\beta,\text{high}})$, and $\theta_{\beta,\text{high}}$. It is equal to $\theta_{\beta,\text{high}}$ if and only if both the fraction of police officers on duty in her precinct of residence and the measure of the strictness of the judge or jury are not less than one half, it is equal to $\theta_{\beta,\text{low}}$ if and only if both of these values are less than one half, otherwise it is equal to $\theta_{\beta,\text{mid}}$. The modified private cost parameters are depicted in the top panels of Figure 2.9; therein, the white area is the set of states of nature on which a private cost parameter is equal to $\theta_{\beta,\text{low}}$, the area shaded in light gray is the set on which it is equal to $\theta_{\beta,\text{mid}}$, and the area shaded in dark gray is the set on which it is equal to $\theta_{\beta,\text{high}}$.

To study the effects of information on equilibrium outcomes, we consider two alternative policy interventions: more private information to a single offender or more public information to all offenders.

First, we consider the case where a single offender has access to more information in the form of an alternative private signal that is more informative than her original private signal. Analogous to the modified private cost parameters, an alternative signal can assume three different values. The four offenders' alternative signals are denoted by $\tilde{s}_{1,\beta}$, $\tilde{s}_{2,\beta}$, $\tilde{s}_{3,\beta}$, $\tilde{s}_{4,\beta}$ and are depicted in the bottom panels of Figure 2.9. Some comments are in order. First, every alternative signal $\tilde{s}_{i,\beta}$ is more informative (about the state of nature) than the original signal $s_{i,\beta}$ because the former induces a finer partition of the set of states nature (the unit square $[0, 1]^2$) than the latter. Second, the alternative signals provide offenders 1, 3, and 4 with complete information about

25. See Section A.5 in the Appendix for a formal discussion.

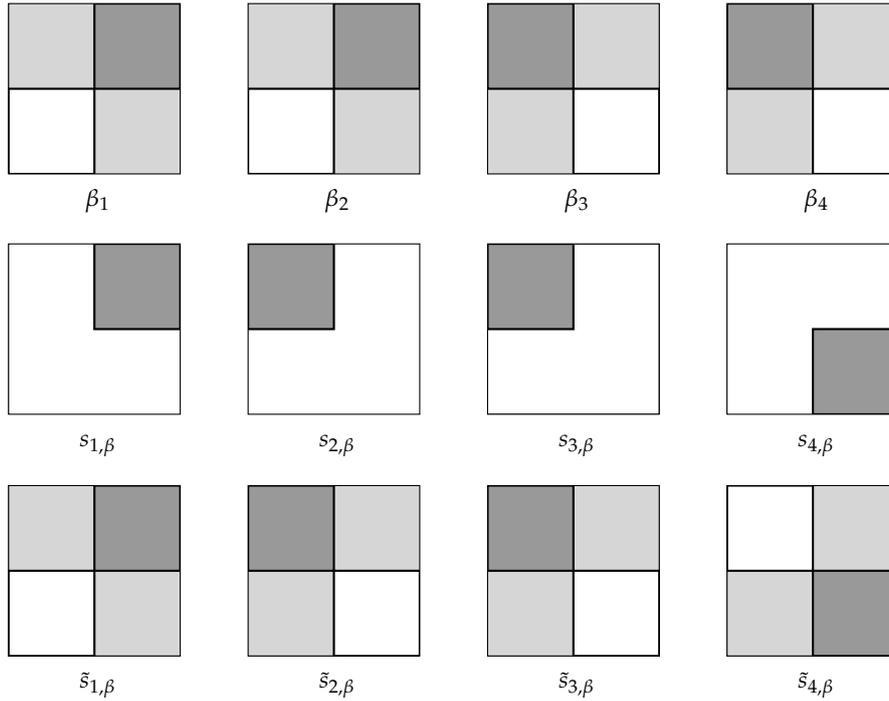


Figure 2.9 Offenders' modified private cost parameters (top panels), original private signals (middle panels) and alternative private signals (bottom panels)

the value of their modified private cost parameters. Third, although the alternative signal $\tilde{s}_{2,\beta}$ induces a finer partition of the set of states of nature than the original signal $s_{2,\beta}$, it does not convey more information to offender 2 about the value of β_2 .²⁶

Second, we consider the case of public information in the form of a public signal all offenders observe in addition to their original private signals. A central planner may, for example, send a public signal in the form of a public announcement regarding the prevention of crime. The public signal can also assume three different values. It is depicted in Figure 2.10. By observing the values of both their original private signals and the public signal, offenders 1, 3, and 4 learn the value of their private cost parameters; in other words, there is no uncertainty about the values of their private cost parameters. The public signal provides, however, no information to offender 2 about the value of her private cost parameter beyond that embodied in her original signal.

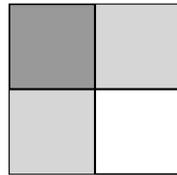


Figure 2.10 Public signal

Table 2.4 reports relative changes in ex ante expected equilibrium efforts and payoffs due to more private information to a single offenders (see columns 2 to 5) or public information

26. Indeed, $\mathbb{E}(\beta_2 | s_{2,\beta}) = \mathbb{E}(\beta_2 | \tilde{s}_{2,\beta}) = \theta_{\beta,\text{mid}}$.

(see the last column), that is, more information to all offenders in the form of a public signal. As can be seen, more private information to a single offender may increase, decrease, or leave unaffected her ex ante expected equilibrium payoff and those of other offenders. A similar result may even hold true for an offender whose payoff does not depend on other offenders' efforts (see Example 6.7). As discussed in Section A.5 in the Appendix, there exist values of the private benefit and social cost parameters such that more private information causes positive, negative, or no externalities on other offenders. The provision of a public signal may decrease ex ante expected equilibrium payoff, as exhibited in the case of offender 2. Therefore, more information is not always beneficial to agents, here offenders.

2.6 Contribution to the crime literature

We believe that our model provides an important contribution to the crime literature.

In the standard model of crime, the probability of apprehension and conviction for a given level of criminal activity is known and common knowledge and, moreover, identical for all individuals. In practice, however, these assumptions may not always be plausible. First, individuals may have imperfect knowledge of the probability of getting apprehended and convicted when deciding whether to commit a crime. Even in the case of a potential conviction, the sanctions may not be certain a priori because they depend on factors such as the exact charges that will be pressed, the leniency of the judge or the jury, and the fact that the penalty for a crime may vary.²⁷ Second, the perception of the probability of apprehension and conviction is not uniform but may vary across individuals or groups of individuals.²⁸

Ultimately, one of the most important questions for perceptual-deterrence research is the degree of correspondence between *actual* and *perceived* risks. In our model, the actual risk (of being caught) of offender i is embodied in β_i while the perceived risk is a predictor of β_i that is based on the information the offender receives from a private signal s_i , more specifically, the conditional expectation of β_i given the signal s_i , (denoted by $\mathbb{E}(\beta_i | s_i)$). Therefore, perceptions will closely mirror reality when the signal is informative about the risk of committing crime and thus the predictor is close to β_i , which means that the variance of the conditional expectation, $\text{var}(\mathbb{E}(\beta_i | s_i))$, is high and close to its maximum value, which is the variance of β_i , $\text{var}(\beta_i)$. In that case, using policy shocks to learn about the magnitude of deterrence is straightforward.

However, if perceptions will not closely mirror reality, which means that the signal is uninformative about β_i so that $\text{var}(\mathbb{E}(\beta_i | s_i))$ is close to zero, then changes in policy will go unnoticed by potential offenders. In that case, the outcomes of policy research will tend to be of limited value in studying deterrence.

In this paper, we have three main contributions to the crime literature. First, we explicitly model the gap between the *actual* and *perceived* risks of committing crime for each offender. For example, Blanes i Vidal and Mastrobuoni (2018) show that, if offenders do not know the presence of police in their area, so that there is a large gap between the *actual* and *perceived* risks of being caught, then increasing police forces in an area has no effect on crime deterrence.²⁹

27. See Bebchuk and Kaplow (1992) on how the introduction of uncertainty affects Becker's conclusions and Polinsky and Shavell (2000) and Lindquist and Zenou (2019) for a survey of various approaches to modelling crime.

28. See Sah (1991) and the references therein.

29. More precisely, Blanes i Vidal and Mastrobuoni (2018) exploit a natural experiment that aimed to increase police presence in more than 6,000 well-defined areas, in Essex, England. Using data transmitted by GPS devices worn by police officers, they do not find that these increases in patrolling were accompanied by corresponding decreases in crime.

Table 2.4 Relative changes in ex ante expected equilibrium efforts and payoffs in the game of crime with network D due to either more private information to a single offender or more information to all offenders in the form of a public signal

Relative changes in ex ante expected equilibrium efforts					
i	$d_r(\mu_{x_i^*}(\beta_1), \mu_{x_i^*}(\beta))$	$d_r(\mu_{x_i^*}(\beta_2), \mu_{x_i^*}(\beta))$	$d_r(\mu_{x_i^*}(\beta_3), \mu_{x_i^*}(\beta))$	$d_r(\mu_{x_i^*}(\beta_4), \mu_{x_i^*}(\beta))$	$d_r(\mu_{x_i^*}(\tilde{\beta}), \mu_{x_i^*}(\beta))$
1	0.0322	0	0.0052	0	0.0370
2	0	0	0.0158	0	0.0179
3	0	0	0.0461	0	0.0523
4	0.0007	0	0.0153	0.0154	0.0358

Relative changes in ex ante expected equilibrium payoffs					
i	$d_r(\mu_{u_i^*}(\beta_1), \mu_{u_i^*}(\beta))$	$d_r(\mu_{u_i^*}(\beta_2), \mu_{u_i^*}(\beta))$	$d_r(\mu_{u_i^*}(\beta_3), \mu_{u_i^*}(\beta))$	$d_r(\mu_{u_i^*}(\beta_4), \mu_{u_i^*}(\beta))$	$d_r(\mu_{u_i^*}(\tilde{\beta}), \mu_{u_i^*}(\beta))$
1	0.0513	0	-0.0009	0	0.0449
2	0	0	-0.0377	0	-0.0342
3	0	0	0.0570	0	0.0702
4	-0.0160	0	0.0187	0.0223	0.0364

Notes: All real numbers in decimal notation are rounded to four decimal places. As regards notation, β denotes the game of crime in which all offenders observe their original private signals, β_i the game in which all but offender i observe their original private signals and offender i her alternative private signal, and $\tilde{\beta}$ the game in which all offenders observe both their original private signal and the public signal. $\mu_{x_i^*}(\beta)$ and $\mu_{u_i^*}(\beta)$ denote offender i 's ex ante expected equilibrium effort and payoff, respectively, in the game of crime β (an analogous notation applies to the games of crime β_i and $\tilde{\beta}$); and the relative change of quantity z relative to the nonzero quantity z_{ref} , $d_r(z, z_{\text{ref}})$ is defined as $(z - z_{\text{ref}}) / |z_{\text{ref}}|$.

Second, we also model the influence of direct friends or co-offenders on own crime activity and, therefore, how the gap between the *actual* and *perceived* risks of committing crime of one’s friends affects own criminal activities.³⁰ Our third contribution is to study policies aiming at reducing crime and to show how the quality of information on the presence of the police and the severity of the judge in case of arrest crucially affects the effectiveness of these policies.³¹

We believe that these are important issues because, in order to address adequate policies aiming at reducing crime, we need to understand how offenders react and change their beliefs to an increase in police and/or an increase in punishment. For example, using Norwegian data, Bhuller et al. (2018) estimate the effect of the punishment of criminals on their co-offenders and their brothers. In fact, using the quasi-random assignment of criminal cases to judges, they show that, when a defendant is sent to prison, there are 51 and 32 percentage point reductions in the probability his criminal network members and younger brothers will be charged with a crime, respectively, over the ensuing four years. Our model can provide an explanation of these empirical results. When a criminal is facing a “tough” judge, she updates her beliefs about the severity of crime punishment and transmit this information to her co-offenders and younger brothers, who, in turn, update their own beliefs, and, as a result, reduce their criminal activities.

3 The Bayesian network game

Let us now derive general results for any number of agents, any information structure and any network. This section specifies the Bayesian network game (Section 3.1) and studies the existence and uniqueness of a Bayesian Nash equilibrium in pure strategies (Section 3.2).

3.1 Specification

The constituent parts of the Bayesian network game are a finite set of players, a common action space, a state space with a common prior, a network connecting the players, the players’ payoff functions, and the players’ signals and type spaces.

Set of players The number of players is finite, at least two, and denoted by I . The set of all players is identified with the set $[I] = \{1, \dots, I\}$.³² All mathematical objects associated with a particular player are indexed by an element of $[I]$, typically by i .

Common action space The players have a common action space, which is equal to \mathbb{R}_+ , the set of nonnegative real numbers. An action of player i is denoted by y_i . The set of all action profiles (y_1, \dots, y_I) is equal to \mathbb{R}_+^I .

30. Indeed, an offender’s equilibrium strategy depends not only on his own predictor but also on the predictors of the offenders in her reference group and higher order reference groups. It can be shown that an offender’s ex ante expected equilibrium effort $\mathbb{E}(x_i^* \circ s_i)$ is a function of $\text{var}(\mathbb{E}(\beta_i \mid s_i))$. See Example E.5, the expression for $\mathbb{E}(x_2^* \circ s_2)$ and Lemma E.1 in particular.

31. See also Cortés, Friebe, and Maldonado (2019) who develop a model in which the information needed to engage in crime activities arrives in the form of a rumour and show that policies that decrease the cost of education for talented students may increase crime participation from less talented students.

32. For any positive integer z , the symbol $[z]$ denotes the set of integers $\{1, \dots, z\}$.

State space and common prior The state space is a measurable space (Ω, \mathfrak{S}) , where the non-empty set Ω represents all possible states of nature that are relevant to the game. With a slight abuse of terminology, the set Ω is also referred to as the state space. The players have a common prior, that is, probability measure, \mathbb{P} on (Ω, \mathfrak{S}) . The probability space $(\Omega, \mathfrak{S}, \mathbb{P})$ constitutes the probabilistic framework of the game.

Network The players are connected by an exogenous network that is constant across the states of nature and common knowledge among the players. The network is represented by a digraph D on $[I]$.^{33,34} The term network is used as a synonym for digraph hereinafter. The network D encodes information about the identities of the players who directly affect a player's payoff through their actions. For a particular player, the set of players who directly affect her payoff corresponds to her *out-neighborhood* in D . Player i 's out-neighborhood in D is denoted by $\mathcal{N}_D^+(i)$ and its cardinality, the so-called *out-degree* of i in D , by $\deg_D^+(i)$. By the definition of a digraph, $i \notin \mathcal{N}_D^+(i)$. All players have at least one out-neighbor in D , that is, for all $i \in [I]$, $\mathcal{N}_D^+(i) \neq \emptyset$ or, equivalently, $\deg_D^+(i) \geq 1$. Note that a player is not necessarily an out-neighbor of her out-neighbors, that is, $n \in \mathcal{N}_D^+(i)$ does not necessarily imply that $i \in \mathcal{N}_D^+(n)$. The dependence of a player's payoff on the actions of her out-neighbors is, therefore, potentially unidirectional.

Payoff functions Player i 's payoff function $u_i: \Omega \times \mathbb{R}_+^I \rightarrow \mathbb{R}$ is defined by:

$$u_i(\omega, (y_1, \dots, y_I)) := \alpha_i(\omega)y_i - \frac{\beta_i(\omega)}{2}y_i^2 - \frac{\gamma_i(\omega)}{2} \left(y_i - \frac{\sum_{n \in \mathcal{N}_D^+(i)} y_n}{\deg_D^+(i)} \right)^2, \quad (2)$$

where $\alpha_i(\omega) > 0$, $\beta_i(\omega) > 0$, $\gamma_i(\omega) \geq 0$ are the values that the square-integrable random variables $\alpha_i: \Omega \rightarrow \mathbb{R}_{++}$, $\beta_i: \Omega \rightarrow \mathbb{R}_{++}$, $\gamma_i: \Omega \rightarrow \mathbb{R}_+$ on $(\Omega, \mathfrak{S}, \mathbb{P})$ assume at the state of nature ω . Note that $\alpha_i, \beta_i, \gamma_i$ are functions, which may in principle depend on some network property, for example, out-degree. The triple $(\alpha_i, \beta_i, \gamma_i)$ is referred to as player i 's *payoff parameters*. The players' payoff parameters $(\alpha_1, \beta_1, \gamma_1), \dots, (\alpha_I, \beta_I, \gamma_I)$ may be functionally or stochastically dependent, and for all $i \in [I]$, $\alpha_i, \beta_i, \gamma_i$ may be functionally or stochastically dependent.³⁵ The common functional form of the payoff functions is common knowledge among the players. Incomplete information may arise by the players' ignorance about the values of some of the payoff parameters. A payoff parameter can give rise to incomplete information only if it is not constant.

The following definition is useful for a compact representation of the payoff functions and the statement of results.

Definition 3.1 The *row-normalized adjacency matrix* of D (with respect to the identity mapping on $[I]$) is the square matrix $\bar{A}(D)$ of order I whose component in row i and column n is defined as $\mathbb{1}_{\mathcal{N}_D^+(i)}(n) / \deg_D^+(i)$ and denoted by $\bar{a}_{i,n}(D)$.³⁶

For all $(i, n) \in [I]^2$, $\bar{a}_{i,n}(D) = 1 / \deg_D^+(i)$ if player n is an out-neighbor of player i in D , and $\bar{a}_{i,n}(D) = 0$ otherwise, and $\bar{a}_{i,i}(D) = 0$.

33. See Appendix E for basic concepts in graph theory.

34. Note that an undirected graph can be represented by a symmetric digraph.

35. For example, the functions $f_1: \Omega \rightarrow \mathbb{R}$, $f_2: \Omega \rightarrow \mathbb{R}$, $f_3: \Omega \rightarrow \mathbb{R}$ are called *functionally dependent* if there exists a nonzero function $g: \mathbb{R}^3 \rightarrow \mathbb{R}$ such that $\omega \mapsto g(f_1(\omega), f_2(\omega), f_3(\omega))$ is identically zero on Ω ; if no such g exists, then f_1, f_2, f_3 are called *functionally independent*.

36. The symbol $\mathbb{1}_{\mathcal{S}}$ denotes the indicator function of the set \mathcal{S} .

Using Definition 3.1, player i 's payoff (2) can be written as

$$u_i(\omega, (y_1, \dots, y_I)) = \alpha_i(\omega)y_i - \frac{\beta_i(\omega)}{2}y_i^2 - \frac{\gamma_i(\omega)}{2}\left(y_i - \sum_{n \in [I]} \bar{a}_{i,n}(D)y_n\right)^2.$$

Some comments on the payoff functions are in order. First, player i 's payoff function is symmetric in her out-neighbors' actions. It exhibits local strategic complements if $\gamma_i > 0$ because for all $n \in [I]$,

$$\frac{\partial^2 u_i(\omega, (y_1, \dots, y_I))}{\partial y_i \partial y_n} = \begin{cases} 0 & \text{if } n \notin \mathcal{N}_D^+(i), \\ \frac{\gamma_i(\omega)}{\deg_D^+(i)} & \text{if } n \in \mathcal{N}_D^+(i). \end{cases}$$

It does, however, not exhibit positive or negative local externalities.³⁷

Second, the payoff that player i ascribes to the action profile (y_1, \dots, y_I) consists of two components: a private component and a social component. The *private component* is defined as $\alpha_i(\omega)y_i - (\beta_i(\omega)/2)y_i^2$, which in turn can be decomposed into two parts: a private benefit and a private cost. The *private benefit* is defined as $\alpha_i(\omega)y_i$ and the *private cost* as $(\beta_i(\omega)/2)y_i^2$. The payoff parameters α_i and β_i are referred to as player i 's *private benefit parameter* and *private cost parameter*, respectively. Note that the marginal private benefit does not vary with y_i and the marginal private cost is strictly increasing in y_i . The *social component* is defined as $(\gamma_i(\omega)/2)(y_i - \sum_{n \in [I]} \bar{a}_{i,n}(D)y_n)^2$; it represents player i 's social cost, if any ($\gamma_i = 0$ is possible), from deviating from a *social norm* that is given by the arithmetic mean of her out-neighbors' actions, $\sum_{n \in [I]} \bar{a}_{i,n}(D)y_n$. The payoff parameter γ_i is referred to as player i 's *social cost parameter*. Note that γ_i is nonnegative, that is, it may be positive for all states of nature, it may be zero for some states of nature, or it may be zero for all states of nature. Also note that the marginal social cost is strictly increasing in y_i and strictly decreasing in $\sum_{n \in [I]} \bar{a}_{i,n}(D)y_n$ if $\gamma_i(\omega) > 0$. The distance between player i 's action and her social norm is referred to as the *social distance* (see also Akerlof 1997) between player i and her out-neighbors. It is important to note that the players' social norms are endogenous and potentially heterogeneous (in equilibrium) because the players may vary in their out-neighborhoods and may choose different actions (in equilibrium).

Third, the payoff functions can be extended to cover the case of players without out-neighbors in D ; specifically, player i 's extended payoff function $\hat{u}_i: \Omega \times \mathbb{R}_+^I \rightarrow \mathbb{R}$ is defined by

$$\hat{u}_i(\omega, (y_1, \dots, y_I)) := \begin{cases} \alpha_i(\omega)y_i - \frac{\beta_i(\omega)}{2}y_i^2 & \text{if } \mathcal{N}_D^+(i) = \emptyset, \\ u_i(\omega, (y_1, \dots, y_I)) & \text{if } \mathcal{N}_D^+(i) \neq \emptyset. \end{cases}$$

The assumption that all players have at least one out-neighbor in D is, in light of this extension, without loss of generality because the extended payoff function of a player without out-neighbors in D corresponds to the original payoff function with a zero social cost parameter.

Signals and type spaces Each player receives information about the unobservable state of nature via a signal whose values are private, that is, not observable by other players. Player i 's

37. In accordance with the terminology introduced by Galeotti et al. (2010, pp. 226–27), player i 's payoff function is said to exhibit negative (respectively, positive) local externalities if for all $\omega \in \Omega$, for all $(y_1, \dots, y_I) \in \mathbb{R}_+^I$, and for all $(\tilde{y}_1, \dots, \tilde{y}_I) \in \mathbb{R}_+^I$ with $\tilde{y}_i = y_i$ and $\{\tilde{y}_n - y_n \mid n \in \mathcal{N}_D^+(i)\} \subset \mathbb{R}_+$, $u_i(\omega, (\tilde{y}_1, \dots, \tilde{y}_I)) \leq u_i(\omega, (y_1, \dots, y_I))$ (respectively, $u_i(\omega, (\tilde{y}_1, \dots, \tilde{y}_I)) \geq u_i(\omega, (y_1, \dots, y_I))$).

signal is a random 3-vector $\mathbf{s}_i := (s_{i,\alpha}, s_{i,\beta}, s_{i,\gamma})$ on the probability space $(\Omega, \mathfrak{G}, \mathbb{P})$ that has a finite support $\Theta_i \subset \mathbb{R}^3$; more precisely, \mathbf{s}_i is a \mathfrak{G} - $\mathcal{B}(\mathbb{R}^3)$ -measurable function from Ω to \mathbb{R}^3 , where $\mathcal{B}(\mathbb{R}^3)$ denotes the σ -field generated by the Euclidean topology on \mathbb{R}^3 . The signal \mathbf{s}_i and its components $s_{i,\alpha}, s_{i,\beta}, s_{i,\gamma}$ can, therefore, be regarded as simple random elements on $(\Omega, \mathfrak{G}, \mathbb{P})$.^{38,39} The set Θ_i corresponds to player i 's type space and is written as $\{\boldsymbol{\theta}_{i,t} \mid t \in [|\Theta_i|]\}$. Note that, by the definition of a support, for all $i \in [I]$ and for all $\boldsymbol{\theta}_i \in \Theta_i$, $\mathbb{P}(\mathbf{s}_i = \boldsymbol{\theta}_i) > 0$. Also note that no assumptions are made about the joint distribution of the signals $\mathbf{s}_1, \dots, \mathbf{s}_I$; more specifically, they may be identically or not identically distributed, and they may be stochastically dependent or independent.

A signal (respectively, signal component) is called *completely informative* about a payoff parameter if the payoff parameter is measurable with respect to the σ -field generated by the signal (respectively, signal component), that is, the payoff parameter is equal to some Borel-measurable function of the signal (respectively, signal component). A signal (respectively, signal component) is called *completely uninformative* about a payoff parameter if the payoff parameter and the signal (respectively, signal component) are stochastically independent.

Timing The timing of the game is as follows:

1. *Nature moves.* Nature determines a state $\omega \in \Omega$, referred to as the *true state of nature*, which is not observed by the players.
2. *Players receive information.* Each player i observes the value of her signal \mathbf{s}_i at the state ω , which determines her type $\boldsymbol{\theta}_i \in \Theta_i$, where $\boldsymbol{\theta}_i = \mathbf{s}_i(\omega)$.
3. *Players move.* Each player i chooses an action $y_i \in \mathbb{R}_+$ conditional on her type $\boldsymbol{\theta}_i$, that is, y_i is player i 's action on the event $\{\mathbf{s}_i = \boldsymbol{\theta}_i\} \subset \Omega$.
4. *Players receive payoffs.* Each player i receives the payoff that corresponds to the state ω and the profile (y_1, \dots, y_I) of actions chosen by all players, $u_i(\omega, (y_1, \dots, y_I))$.

The Bayesian network game specified above is denoted by \mathcal{B} . The constituent parts of \mathcal{B} , which include, inter alia, the prior \mathbb{P} , the signals $(\mathbf{s}_i)_{i \in [I]}$, the network D , and the payoff parameters $((\alpha_i, \beta_i, \gamma_i))_{i \in [I]}$, are collectively referred to as the *structure* of \mathcal{B} . Note that \mathcal{B} is a game with complete information if there is only one state of nature, that is, $|\Omega| = 1$, in which case all signals and all preference parameters are constant.

3.2 Existence and uniqueness of Bayesian Nash equilibrium

A pure strategy of player i is a function $x_i: \Theta_i \rightarrow \mathbb{R}_+$, that is, a rule that assigns a unique action to each type. The composition of a pure strategy x_i with \mathbf{s}_i , $x_i \circ \mathbf{s}_i: \Omega \rightarrow \mathbb{R}_+$ is a random variable on $(\Omega, \mathfrak{G}, \mathbb{P})$. The set of all pure strategies of player i is denoted by $\mathbb{R}_+^{\Theta_i}$. The set of all pure strategy profiles is equal to $\times_{i \in [I]} \mathbb{R}_+^{\Theta_i}$.

The analysis of existence and uniqueness of equilibrium is relegated to Appendix B. First, we show that there exists a unique interior Bayesian-Nash equilibrium (BNE) in Proposition B.1 and characterize the equilibrium strategies. Then, Proposition B.2 gives sufficient conditions for

38. A random element is called simple if it assumes a finite number of values.

39. We do not distinguish between random elements that are equal almost surely.

a symmetric BNE that is constant across the states of nature and Corollary B.2 determines the equilibrium strategies when all signals are constant.⁴⁰

4 Comparative statics analysis

This section studies how equilibrium behavior, ex ante expected equilibrium strategies in particular, hereinafter referred to as ex ante expected equilibrium actions, responds to changes in payoff parameters (Section 4.1) and changes in the network (Section 4.2).

A change in a player's payoff parameter or out-neighborhood may affect not only her ex ante expected equilibrium action but also those of her in-neighbors and higher-order in-neighbors. The effect of such a change propagates within the network along the inverses of walks from the player's in-neighbors or higher-order in-neighbors to the player only if the walks have certain properties pertaining to the structure of the Bayesian network game and the nature of the change. Definitions 4.1, 4.2, and 4.3 introduce the relevant terminology. Example C.1 in Appendix C illustrates Definitions 4.2 and 4.3.

Definition 4.1 A player is called *quasi-isolated* if her social cost parameter is zero a.s.

A quasi-isolated player is not affected by the behavior of her out-neighbors but may affect the behavior of existing in-neighbors; specifically, if player i is quasi-isolated, then her equilibrium strategy and ex post expected equilibrium payoff $u_i^* := u_i(\text{id}_\Omega, (x_1^* \circ s_1, \dots, x_1^* \circ s_1))$ are given by

$$x_i^* \circ s_i = \frac{\mathbb{E}(\alpha_i | s_i)}{\mathbb{E}(\beta_i | s_i)} \quad \text{and} \quad \mathbb{E}(u_i^* | s_i) = \frac{1}{2} \frac{\mathbb{E}(\alpha_i | s_i)^2}{\mathbb{E}(\beta_i | s_i)}. \quad (3)$$

Definition 4.2 A walk (i_0, \dots, i_p) in D of length p from player i_0 to player i_p is called *conductive* (in the Bayesian network game \mathcal{B}) if there exists a $(p+1)$ -tuple $(\vartheta_{i_0}, \dots, \vartheta_{i_p})$ in $\times_{z=0}^p \Theta_{i_z}$ such that for all $z \in [p]$, $\mathbb{P}(\gamma_{i_{z-1}} > 0, s_{i_{z-1}} = \vartheta_{i_{z-1}}, s_{i_z} = \vartheta_{i_z}) > 0$, in which case $\{s_{i_p} = \vartheta_{i_p}\}$ is called a *head event* of the walk (i_0, \dots, i_p) .

Provided that all signals are constant, a walk (i_0, \dots, i_p) in D is conductive if and only if players i_0, \dots, i_{p-1} are not quasi-isolated. Provided that not all signals are constant, a walk (i_0, \dots, i_p) in D is conductive if there exists a $(p+1)$ -tuple $(\vartheta_{i_0}, \dots, \vartheta_{i_p})$ in $\times_{z=0}^p \Theta_{i_z}$ such that for all $z \in [p]$, $\mathbb{P}(\gamma_{i_{z-1}} > 0) = 1$ and $\mathbb{P}(s_{i_{z-1}} = \vartheta_{i_{z-1}}, s_{i_z} = \vartheta_{i_z}) > 0$.⁴¹

Definition 4.3 A conductive walk in D is said to have a *positive* (respectively, *negative*) *intersection* with a random variable π on $(\Omega, \mathfrak{S}, \mathbb{P})$ if the walk has a head event for which the intersection with the event $\{\pi > 0\}$ (respectively, $\{\pi < 0\}$) has positive probability.

Provided that all signals are constant, a conductive walk in D has a positive (respectively, negative) intersection with a random variable π on $(\Omega, \mathfrak{S}, \mathbb{P})$ if and only if π is positive (respectively, negative) with positive probability.

For future reference, we introduce two independence conditions involving the private or social cost parameters and the signals (Conditions 4.4 (1) and 4.4 (2)) and a condition stating that changes in the topology of the network do not affect the payoff parameters (Condition 4.5).

40. Table 2.1 in Section 2 provides the ex ante equilibrium efforts, payoffs and social norms for our criminal network example with 4 players.

41. For all $z \in [p]$, $\mathbb{P}(s_{i_{z-1}} = \vartheta_{i_{z-1}}, s_{i_z} = \vartheta_{i_z}) = \mathbb{P}(\gamma_{i_{z-1}} > 0, s_{i_{z-1}} = \vartheta_{i_{z-1}}, s_{i_z} = \vartheta_{i_z})$ if $\mathbb{P}(\gamma_{i_{z-1}} > 0) = 1$.

Condition 4.4 (1) For all $i \in [I]$, β_i and s_i are stochastically independent.

(2) For all $i \in [I]$ and for all $n \in \{i\} \cup \mathcal{N}_D^+(i)$, γ_i and s_n are stochastically independent.

Condition 4.4 (1) states that a player's signal is completely uninformative about her private cost parameter, in which case the best (in terms of mean squared error) predictor of the private cost parameter that is a function of the signal is its expectation. Condition 4.4 (2) states that a player's signal is completely uninformative about her social cost parameter, and the same applies to the signals of the player's out-neighbors. Condition 4.4 is satisfied if all signals are constant across the states of nature or all private and social cost parameters are constant across the states of nature; the condition is, therefore, satisfied, if there is only one state of nature, which covers the case where the players have complete information.

Condition 4.5 For all $i \in [I]$, α_i , β_i , and γ_i do not vary with the arc set of D .

4.1 Changes in payoff parameters

The analysis allows for two types of changes in payoff parameters: global and individual. The payoff parameters are to this end decomposed into two summands: for all $i \in [I]$,

$$\alpha_i = \alpha^G + \alpha_i^L, \quad \beta_i = \beta^G + \beta_i^L, \quad \gamma_i = \gamma^G + \gamma_i^L,$$

where α^G is referred to as the *global component* of α_i , which is common to all players and, for example, defined as $(1/I) \sum_{n \in [I]} \alpha_n$ or $\min\{\alpha_n \mid n \in [I]\}$, and $\alpha_i^L := \alpha_i - \alpha^G$ is referred to as the *idiosyncratic component* of α_i . The parameter components β^G , β_i^L , γ^G , γ_i^L are defined analogously.

The global component of a payoff parameter characterizes some attribute of the activity under consideration that does not directly depend on an individual's characteristics, whereas these are represented by the idiosyncratic component. Consider, for example, optimal investment in education. The global component α^G corresponds to the average marginal increase in earnings from an additional year of schooling, and the global component β^G captures the additional average cost (for example, tuition fees or foregone income). A student may possess skills that enables her to benefit more than the average student from an additional year of schooling, in which case her idiosyncratic component α_i^L is positive, while the opposite may be true for a less skilled student, in which case his α_i^L is negative. Costs may vary among students as well. A high-ability student is more likely to receive a scholarship than a low-ability student, in which case the idiosyncratic component β_i^L of the high-ability student is negative. The opportunity cost of obtaining a postgraduate degree may be higher for an individual who is already employed compared to an individual who is unemployed or has just graduated from college, in which case the postgraduate student has a positive β_i^L . The global component γ^G represents the average strength of the prevailing social norm in society. While some individuals may feel more compelled to adhere to the social norm, which corresponds to a positive idiosyncratic component γ_i^L , others may feel less compelled to do so, in which case γ_i^L is negative.

The analysis is confined to a specific type of change in a parameter component, namely, a first-order stochastic dominance (FOSD for short) upward shift. Formally, such a shift is modelled by a nonnegative random variable on $(\Omega, \mathfrak{S}, \mathbb{P})$ that is positive with positive probability. For example, a FOSD upward shift $\Delta\alpha^L$ in the idiosyncratic component of player k 's private benefit parameter changes her private benefit parameter from α_k to $\alpha_k + \Delta\alpha^L$ and leaves the private

benefit parameters of all other players unchanged; and a FOSD upward shift $\Delta\alpha^G$ in the global component of the private benefit parameters changes, for all $i \in [I]$, player i 's private benefit parameter from α_i to $\alpha_i + \Delta\alpha^G$.

The comparative statics results are grouped into two propositions. Results about the signs of the effects of FOSD upward shifts in idiosyncratic and global parameter components on ex ante expected equilibrium actions are stated in Proposition 4.6. Results about the variation of the signs of these effects among affected players are stated in Proposition 4.7.

Proposition 4.6 *Let $(j, k) \in [I]^2$ with $j \neq k$.*

- (1) *A FOSD upward shift $\Delta\alpha^L$ in player k 's idiosyncratic component of the private benefit parameter strictly increases her ex ante expected equilibrium action; it strictly increases player j 's ex ante expected equilibrium action if and only if there exists a conductive walk in D from player j to player k that has a positive intersection with $\Delta\alpha^L$.*
- (2) *A FOSD upward shift in the global component of the private benefit parameters strictly increases all ex ante expected equilibrium actions.*
- (3) *A FOSD upward shift $\Delta\beta^L$ in player k 's idiosyncratic component of the private cost parameter strictly decreases her ex ante expected equilibrium action; it strictly decreases player j 's ex ante expected equilibrium action if and only if there exists a conductive walk in D from player j to player k that has a positive intersection with $\Delta\beta^L$.*
- (4) *A FOSD upward shift in the global component of the private cost parameters strictly decreases all ex ante expected equilibrium actions.*
- (5) *A FOSD upward shift in a player's idiosyncratic component of the social cost parameter may decrease, increase, or leave unchanged her and other players' ex ante expected equilibrium actions; specifically, a FOSD upward shift $\Delta\gamma^L$ in player k 's idiosyncratic component of the social cost parameter*
 - (a) *strictly decreases (respectively, increases) her ex ante expected equilibrium action if her equilibrium strategy $x_k^* \circ s_k$ is greater (respectively, less) than her social norm $\sum_{n \in [I]} \bar{a}_{k,n}(D)(x_n^* \circ s_n)$ on at least one event on which $\Delta\gamma^L$ is positive with positive probability and greater than or equal to (respectively, less than or equal to) her social norm on all other events on which $\Delta\gamma^L$ is positive with positive probability;*
 - (b) *strictly decreases (respectively, increases) player j 's ex ante expected equilibrium action if player k 's strategy is greater than or equal to (respectively, less than or equal to) her social norm on all events on which $\Delta\gamma^L$ is positive with positive probability and there exists a conductive walk in D from player j to player k that has a positive (respectively, negative) intersection with $\Delta\gamma^L(x_k^* \circ s_k - \sum_{n \in [I]} \bar{a}_{k,n}(D)(x_n^* \circ s_n))$;*
 - (c) *does not change player j 's ex ante expected equilibrium action if there does not exist a conductive walk in D from player j to player k that has a positive intersection with $\Delta\gamma^L$.*
- (6) *A FOSD upward shift in the global component of the social cost parameters may decrease, increase, or leave unchanged a player's ex ante expected equilibrium action. A FOSD upward shift $\Delta\gamma^G$ in the global component of the social cost parameters strictly decreases (respectively, increases) player i 's ex ante expected equilibrium action if all equilibrium strategies are greater than or equal to (respectively, less than or equal to) their social norms on all events on which $\Delta\gamma^G$ is positive with positive probability*

and player i 's equilibrium strategy is greater (respectively, less) than her social norm on at least one event on which $\Delta\gamma^G$ is positive with positive probability.

For the discussion of Proposition 4.6, suppose there is one state of nature (which corresponds to the case of complete information, where strategies are the same as ex ante expected actions), player k is an out-neighbor of player j , player k is quasi-isolated, and player j is not quasi-isolated. Evidently, (j, k) is a conductive walk in D from player j to player k .

First, consider an upward shift in the idiosyncratic component of player k 's private benefit parameter. The shift creates an imbalance between player k 's marginal private benefit, which does not vary with action, and her marginal private cost, which increases linearly with action; specifically, marginal private benefit exceeds marginal private cost. To correct this imbalance, player k increases her action and therefore her marginal private cost until it equals her marginal private benefit. The increase in player k 's action increases player j 's social norm, which creates an imbalance between his marginal private benefit and his marginal cost, the sum of private and social marginal cost, both of which increase linearly with action. The increase in player j 's social norm decreases his marginal social cost, so that his marginal cost falls short of his marginal private benefit. To correct this imbalance, he increases his action.

Second, consider an upward shift in the global component of the private benefit parameters. As a consequence, all marginal private benefits exceed their marginal costs. To correct this imbalance, all players increase their actions.

Third, consider an upward shift in the idiosyncratic component of player k 's private cost parameter. As a result, player k 's marginal private cost, which increases linearly with action, falls short of her marginal private benefit. To correct this imbalance, player k decreases her action. The decrease in player k 's action decreases player j 's social norm, which creates an imbalance between his marginal private benefit and marginal cost; specifically, marginal private benefit falls short of marginal cost. To correct this imbalance, player j decreases his action.

Fourth, consider an upward shift in the global component of the private cost parameters. As a consequence, all marginal private benefits fall short of their marginal costs. To correct this imbalance, all players decrease their actions.

Fifth, consider an upward shift in the idiosyncratic component of player k 's social cost parameter, so that she is no longer quasi-isolated. As a consequence, player k 's marginal social cost increases (respectively, decreases) if her action is greater (respectively, less) than her social norm. To correct the imbalance between her marginal private benefit and marginal cost, she decreases (respectively, increases) her action. The decrease (respectively, increase) in player k 's action decreases (respectively, increases) player j 's social norm. To correct the resulting imbalance between his marginal private benefit and his marginal cost, player j decreases (respectively, increases) his action.

Sixth, consider an upward shift in the global component of the social cost parameters. As a consequence, the marginal social costs increase (respectively, decrease) for those players whose actions are greater (respectively, less) than their social norms. If the signs of the changes in the marginal social costs vary among players, then the signs of the resulting effects on their actions may vary also.

Results 5 and 6 of Proposition 4.6 assert that the signs of the effects depend on the structure of the Bayesian network game; Example C.1 (see Appendix C) elaborates on this dependence. Moreover, the signs of the effects may vary among affected players; Example C.2 (see Appendix C) and the example in Section A.2 (see Appendix A) discuss two such cases. The latter example

shows in addition that the magnitudes of the effects may not decrease with the distance to the player whose idiosyncratic component of the social cost parameter changes, more specifically, the magnitudes of the effects may not decrease monotonically along the inverses of walks connecting players; in other words, there is in general no effect attenuation within the network. Proposition 4.7 provides a sufficient condition for the signs to be the same among affected players in case of a FOSD upward shift in a player's idiosyncratic component of the social cost parameter.

Proposition 4.7 *Let $\Delta\gamma^L$ be a FOSD upward shift in player k 's idiosyncratic component of the social cost parameter. Suppose Condition 4.4 is satisfied in the Bayesian network games without and with the shift $\Delta\gamma^L$. The signs of the effects of the shift $\Delta\gamma^L$ on ex ante expected equilibrium actions are the same for all affected players. If a player other than player k is affected by the shift, then the magnitude of the effect on her ex ante expected equilibrium action is less than that of player k .*

Proposition 4.7 highlights an important difference between the cases of complete and incomplete information: while the signs of the effects of a FOSD upward shift in the idiosyncratic component of a social cost parameter on ex ante expected equilibrium actions are the same for all affected players in case of complete information, they may vary among affected players in case of incomplete information. Note that the signs of the effects of a FOSD upward shift in the global component of the social cost parameters on ex ante expected equilibrium actions may vary among affected players in case of complete information.

4.2 Changes in the network

The analysis is confined to a particular change in the network: the addition of a single arc, that is, the addition of a new out-neighbor to a player's out-neighborhood.

A new arc from player k to player l may decrease, increase, or leave unchanged player k 's and other players' ex ante expected equilibrium actions. The signs of the effects depend on the structure of the Bayesian network game; they are the same for all players who are affected by the new arc under complete information and may vary among affected players under incomplete information. These results are stated in more detail in Propositions 4.8 and 4.9.

Proposition 4.8 *Let $(j, k, l) \in [I]^3$ with $j \neq k$ and $k \neq l$. Suppose there is no arc in D from player k to player l , player k is not quasi-isolated, and Condition 4.5 is satisfied. A new arc from player k to player l*

- (1) *strictly decreases (respectively, increases) player k 's ex ante expected equilibrium action if player l 's strategy $x_l^* \circ s_l$ is less (respectively, greater) than player k 's social norm $\sum_{n \in [I]} \bar{a}_{k,n}(D)(x_n^* \circ s_n)$ on at least one event on which player k 's social cost parameter γ_k is positive with positive probability and less than or equal to (respectively, greater than or equal to) her social norm on all other events on which her social cost parameter is positive with positive probability;*
- (2) *strictly decreases (respectively, increases) player j 's ex ante expected equilibrium action if player l 's strategy $x_l^* \circ s_l$ is less than or equal to (respectively, greater than or equal to) player k 's social norm $\sum_{n \in [I]} \bar{a}_{k,n}(D)(x_n^* \circ s_n)$ on all events on which player k 's social cost parameter γ_k is positive with positive probability and there exists a conductive walk in D from player j to player k that has a negative (respectively, positive) intersection with $\gamma_k(x_l^* \circ s_l - \sum_{n \in [I]} \bar{a}_{k,n}(D)(x_n^* \circ s_n))$.*

For the discussion of Proposition 4.8, suppose all signals are constant and player k is not quasi-isolated. Let $D + (k, l)$ denote the network that results from D by adding a new arc from player k to

player l . Note that $x_l^* \circ s_l \geq \sum_{n \in [I]} \bar{a}_{k,n}(D)(x_n^* \circ s_n)$ if and only if $\sum_{n \in [I]} \bar{a}_{k,n}(D + (k, l))(x_n^* \circ s_n) \geq \sum_{n \in [I]} \bar{a}_{k,n}(D)(x_n^* \circ s_n)$.⁴² A new arc from player k to player l causes a change in player k 's equilibrium strategy that is of the same sign as the difference between player l 's strategy $x_l^* \circ s_l$ (the minuend) and player k 's social norm $\sum_{n \in [I]} \bar{a}_{k,n}(D)(x_n^* \circ s_n)$ (the subtrahend). Specifically, player k 's equilibrium strategy decreases (respectively, increases) if her new out-neighbor's equilibrium strategy is less (respectively, greater) than the average equilibrium strategies of her current out-neighbors, that is, her out-neighbors in D ; in other words, player k 's equilibrium strategy decreases (respectively, increases) if her social norm decreases (respectively, increases) with the addition of the new out-neighbor, that is, $\sum_{n \in [I]} \bar{a}_{k,n}(D + (k, l))(x_n^* \circ s_n) < \sum_{n \in [I]} \bar{a}_{k,n}(D)(x_n^* \circ s_n)$.⁴³ Apart from this direct effect, the new arc may also have an indirect effect on player k 's in-neighbors and higher-order in-neighbors as it may affect their social norms and in turn their equilibrium strategies. The indirect effect on other players' strategies propagates within the network D along the inverses of walks connecting players if they are not quasi-isolated and player l 's strategy is different from player k 's social norm.

In the context of criminal networks, if a criminal expands her professional network towards criminals who are criminally more (respectively, less) active than the average criminal in her current professional network, then her criminal activity increases (respectively, decreases).

It is interesting to note that a player who is affected by the new arc does not necessarily conform more to the average behavior of her out-neighbors—in the sense that the social distance between the player and her out-neighbors decreases—than before the addition of the arc.

Proposition 4.9 *Let $(k, l) \in [I]^2$ with $k \neq l$. Suppose there is no arc in D from player k to player l and Conditions 4.4 and 4.5 are satisfied in $D + (k, l)$. The signs of the effects of a new arc from player k to player l on ex ante expected equilibrium actions are the same for all players who are affected by the new arc. If a player other than player k is affected by the new arc, then the magnitude of the effect on her ex ante expected equilibrium action is less than that of player k .*

Similar to Proposition 4.7, Proposition 4.9 highlights a difference between the cases of complete and incomplete information: while the signs of the effects on ex ante expected equilibrium actions are the same for all players who are affected by the new arc under complete information, they may vary among affected players under incomplete information. The example in Section A.3 (see Appendix A) illustrates the case where the signs of the effects vary among affected players; in addition, it shows that the magnitude of the effect of a new arc on ex ante expected equilibrium actions is in general not monotonically decreasing along the inverses of walks connecting players.

Propositions 4.7 and 4.9 suggest that the variation of ex ante expected equilibrium actions caused by exogenous shocks to social cost parameters or to the topology of the network connecting the players is not larger under incomplete information than under complete information. Less information may therefore dampen the impact of exogenous shocks at the macro level in terms of a smaller variation in ex ante expected aggregate equilibrium actions.

42. This equivalence follows from the identity

$$\sum_{n \in [I]} \bar{a}_{k,n}(D + (k, l))(x_n^* \circ s_n) - \sum_{n \in [I]} \bar{a}_{k,n}(D)(x_n^* \circ s_n) = \frac{1}{\deg_D^+(k) + 1} \left(x_l^* \circ s_l - \sum_{n \in [I]} \bar{a}_{k,n}(D)(x_n^* \circ s_n) \right).$$

43. Note that both social norms $\sum_{n \in [I]} \bar{a}_{k,n}(D + (k, l))(x_n^* \circ s_n)$ and $\sum_{n \in [I]} \bar{a}_{k,n}(D)(x_n^* \circ s_n)$ are with respect to the BNE (x_1^*, \dots, x_I^*) in the Bayesian network game \mathcal{B} where the players are connected by the network D .

5 Policy analysis

Consider a central planner who knows the structure of the Bayesian network game \mathcal{B} and whose sole objective is to decrease to a degree ex ante expected aggregate equilibrium actions (hereinafter referred to as ex ante expected aggregate action) as, for example, in the context of crime.

The discussion is structured as follows. Section 5.1 assess the effectiveness of policy instruments to decrease ex ante expected aggregate action. Section 5.2 studies optimal targeting by comparing two alternative policies: targeting a single player and targeting all players.

5.1 Effective policy instruments

A policy instrument is called *effective* if there exists a monotone relation between the instrument and ex ante expected aggregate action regardless of the structure of \mathcal{B} . The payoff parameters and the size or density of the network connecting the players are potential candidates for effective policy instruments.

Payoff parameters

Regardless of the structure of \mathcal{B} , ex ante expected aggregate action is strictly increasing in the idiosyncratic component of a player's private benefit parameter and the global component of the private benefit parameters (Results 1 and 2 of Proposition 4.6) and strictly decreasing in the idiosyncratic component of a player's private cost parameter and the global component of the private cost parameters (Results 3 and 4 of Proposition 4.6). Ex ante expected aggregate action is in general not monotone in the idiosyncratic component of a player's social cost parameter and the global component of the social cost parameters (Results 5 and 6 of Proposition 4.6); Example C.1 (see Appendix C) serves as an illustration. These results show that the private benefit parameters and the private cost parameters are effective policy instruments; the social cost parameters are, however, not effective policy instruments.

In the context of education, where action is some measure of educational effort like hours of study time or years of schooling, a policy measure to increase ex ante expected average or aggregate educational effort consists of raising the private benefits or lowering the private costs of educational effort; for example, the subsidization of higher institutions corresponds to a downward shift in the global component of the private cost parameters, scholarships to low-income students correspond to a downward shift in the idiosyncratic components of their private cost parameters, and a policy measure to offsets a decline in income during an economic recession corresponds, for example, to a downward shift in the idiosyncratic components of the private cost parameters of low-income students, where the changes in the idiosyncratic components are negative for all states of nature that are associated with an economic recession and zero for those associated with economic prosperity or growth.

In the context of crime, a policy measure to decrease criminal activity consists of lowering the private benefits or raising the private costs of criminal activity. The theory of deterrence posits that the ex ante expected costs of criminal activity increase with the certainty of punishment (which depends on the probabilities of detection, apprehension, and conviction) and the severity of punishment (as measured, for example, by the level of fines or the length of prison sentences). Consider a network of offenders that spreads across several geographical districts (for example, police areas) of a jurisdiction. A policy measure that increases the certainty of punishment in

all districts corresponds to an upward shift in the global component of the offenders' private cost parameters, and a measure that increases the certainty in a single district corresponds to an upward shift in the idiosyncratic component of only those offenders who operate within that district.

Size or density of network

Ex ante expected aggregate action is in general not monotone in the size or density of the network D (Proposition 4.8); the Example in Section A.3 (see Appendix A), the graph in the left panel of Figure A.6 in particular, demonstrates that the size of D is not an effective policy instrument.⁴⁴ A policy measure to decrease the size of D could, for example, consist of isolating a player (or a group of players) from all other players by severing all her (or their) connections to all other players. In the context of crime, for example, a convicted offender is isolated from her network of criminal associates if the sentence or probation conditions involve association, location, or residence restrictions.

5.2 Optimal targeting

We study optimal targeting by comparing two alternative policies that are defined with respect to either of two effective policy instruments: the private benefit parameters or the private cost parameters.

Policies with respect to private benefit parameters

The two defining elements of a policy with respect to the private benefit parameters are: a set of players called targets and a FOSD downward shift in the private benefit parameters of the targets. The set of targets consists either of a single player, in which case the policy is called a *key player policy*, or all players, in which case the policy is called a *global policy*. A FOSD downward shift is modelled by a nonpositive random variable on the probability space $(\Omega, \mathfrak{S}, \mathbb{P})$ that is negative with positive probability. A FOSD downward shift $\Delta\alpha$ is called α -admissible (in \mathcal{B}) if for all $i \in [I]$, $\alpha_i + \Delta\alpha > 0$. Both elements defining a policy are publicly announced by the central planner and, therefore, common knowledge among the players.

A key player policy (with respect to the private benefit parameters) is defined with reference to the following key player problem (KPP for short), which is abbreviated to KPP- α .⁴⁵

Definition 5.1 (KPP- α) Find the player (in \mathcal{B}) for which a given α -admissible FOSD downward shift in her private benefit parameter results in the maximal decrease of ex ante expected aggregate action.

Let KPP- $\alpha(\Delta\alpha)$ denote the KPP- α with α -admissible FOSD downward shift $\Delta\alpha$. Note that its solution set, that is, the players who solve the KPP- $\alpha(\Delta\alpha)$, is not empty but not necessarily a singleton. A player in the solution set of the KPP- $\alpha(\Delta\alpha)$ is called a *key player* of the KPP- $\alpha(\Delta\alpha)$.

44. The density of a network of order I is defined as the ratio of its size, that is, the number of its arcs, to the maximum number of its arcs, $I(I - 1)$.

45. The selection or identification of a player (or a group of players) that minimizes or maximizes a certain objective function is an important aspect in network analysis, which is subsumed under the term key player (or key group) analysis. See Zenou (2016) for a comprehensive survey of the economics literature on the identification of key players in networks.

A key player policy (with respect to the private benefit parameters) consists of an α -admissible FOSD downward shift $\Delta\alpha^L$ in the private benefit parameter of a single key player of the KPP- $\alpha(\Delta\alpha^L)$. A global policy (with respect to the private benefit parameters) consists of an α -admissible FOSD downward shift $\Delta\alpha^G$ in all private benefit parameters. A key player policy with α -admissible FOSD downward shift $\Delta\alpha^L$ and a global policy with α -admissible FOSD downward shift $\Delta\alpha^G$ are called *comparable* if they satisfy the equality $\mathbb{E}(\Delta\alpha^G) = (1/I) \mathbb{E}(\Delta\alpha^L)$, which can be regarded as representing the central planner's binding budget constraint. One of two such comparable policies is called *weakly superior* (respectively, *strictly superior*) to the other if it decreases ex ante expected aggregate action not less (respectively, more) than the other. Note that the shifts $\Delta\alpha^L$ and $\Delta\alpha^G$ of two comparable policies need not satisfy $\Delta\alpha^G = (1/I)\Delta\alpha^L$; for example, $\Delta\alpha^L$ may be negative at a state of nature at which $\Delta\alpha^G$ is zero.

A key player policy (with respect to the private benefit parameters) is in general neither weakly nor strictly superior to a comparable global policy (with respect to the private benefit parameters), and vice versa. Proposition 5.2 states sufficient conditions for a key player policy to be weakly or strictly superior to a comparable global policy.

Proposition 5.2 *Suppose $\Delta\alpha^G = (1/I)\Delta\alpha^L$ a.s. or all signals are completely uninformative about $\Delta\alpha^L$ and $\Delta\alpha^G$. A key player policy with α -admissible FOSD downward shift $\Delta\alpha^L$ is weakly superior to a comparable global policy with α -admissible FOSD downward shift $\Delta\alpha^G$. If the number of key players of the KPP- $\alpha(\Delta\alpha^L)$ is less than I , then a key player policy is strictly superior to a comparable global policy.*

Targeting a key player is weakly superior to targeting all players in case of complete information (which is covered by the case of uninformative signals), irrespective of how the shifts in the policy instruments are designed by the central planner—their magnitudes may, for example, depend on the economic cycle in different ways.⁴⁶ Targeting a key player is weakly superior to targeting all players in case of incomplete information if the shifts are constant a.s. (which is covered by the case where $\Delta\alpha^G = (1/I)\Delta\alpha^L$ a.s.).⁴⁷

Policies with respect to private cost parameters

A key player policy (with respect to the private cost parameters) is defined with reference to the following key player problem, which is abbreviated to KPP- β .

Definition 5.3 (KPP- β) Find the player (in \mathcal{B}) for which a given FOSD upward shift in her private cost parameter results in the maximal decrease of ex ante expected aggregate action.

A key player policy (with respect to the private cost parameters) consists of a FOSD upward shift $\Delta\beta^L$ in the private cost parameter of a single key player of the related KPP- β . A global policy (with respect to the private cost parameters) consists of a FOSD upward shift $\Delta\beta^G$ in all private cost parameters. A key player policy with FOSD upward shift $\Delta\beta^L$ and a global policy with FOSD upward shift $\Delta\beta^G$ are called *comparable* if they satisfy the equality $\mathbb{E}(\Delta\beta^G) = (1/I) \mathbb{E}(\Delta\beta^L)$.

A key player policy with FOSD upward shift $\Delta\beta^L$ is in general neither weakly nor strictly superior to a comparable global policy with FOSD upward shift $\Delta\beta^G$, and vice versa, even if

46. If there is only one state of nature, then all signals are constant, in which case they are completely uninformative about $\Delta\alpha^L$ and $\Delta\alpha^G$.

47. If $\Delta\alpha^L$ and $\Delta\alpha^G$ are constant a.s., then $\mathbb{E}(\Delta\alpha^G) = (1/I) \mathbb{E}(\Delta\alpha^L)$ (which is satisfied by comparable key player and global policies) is equivalent to $\Delta\alpha^G = (1/I)\Delta\alpha^L$ a.s.

$\Delta\beta^G = (1/I)\Delta\beta^L$ a.s. or all signals are completely uninformative about $\Delta\beta^L$ and $\Delta\beta^G$ (cf. Proposition 5.2). The superiority of one policy over the other depends on the two shifts in the policy instrument, their magnitudes in particular, and the structure of the Bayesian network game \mathcal{B} ; the Example in Section A.4 (see Appendix A) serves as an illustration. In the context of crime where the policy instrument is, for example, the probability or severity of punishment, a key player policy may be strictly superior to a comparable global policy if and only if the shifts in the policy instrument are below certain thresholds.

6 The value of information

A discussion of the value of information must answer two questions; namely, how information is defined and how its value is measured. Section 6.1 addresses the first question. Apart from defining information, it also introduces a measure of information, referred to as informativeness, which proves useful in the statement of formal results. Section 6.2 addresses the second question and details on the relation between the measure of information and the quantity representing the value of information. Drawing on the notation and vocabulary introduced in Section 6.1 and the results presented in Section 6.2, Section 6.3 discusses the value of private information and Section 6.4 the value of public information.

6.1 Informativeness of signal about payoff parameter

A player's signal induces a partition of the state space, which represents her information about the true state of nature (Radner 1968; Marschak and Radner 1972). The finer the partition of the state space, the more information is carried by the signal inducing the partition. For countable partitions, a finer partition generates a larger (in terms of set inclusion) σ -field, and a larger σ -field comes from a finer partition (Hervés-Beloso and Klinger Monteiro 2013, Propositions 1 and 4). A player's information can, therefore, be equivalently represented by the partition induced by her signal or the σ -field generated by her signal.⁴⁸ In the following, the informational content of a signal is represented by the σ -field it generates. The higher the informational content of a signal in terms of the largeness of the σ -field it generates, the greater the variability of the conditional expectation of a payoff parameter given the signal. A natural measure for the informational content of a signal about a payoff parameter is, therefore, the variance of the conditional expectation of the payoff parameter given the signal, which is referred to as the informativeness of the signal about the payoff parameter.

As an illustration, consider player i and the payoff parameter α_i . Suppose the private benefit parameter α_i is not constant, for example, suppose it assumes three different values with equal probability, $\theta_{\alpha,\text{low}}$, $\theta_{\alpha,\text{high}}$, and $\theta_{\alpha,\text{mid}} := (1/2)(\theta_{\alpha,\text{low}} + \theta_{\alpha,\text{high}})$, so that

$$\mathbb{E}(\alpha_i) = \theta_{\alpha,\text{mid}} \quad \text{and} \quad \text{var}(\alpha_i) = \frac{(\theta_{\alpha,\text{low}} - \theta_{\alpha,\text{mid}})^2 + (\theta_{\alpha,\text{high}} - \theta_{\alpha,\text{mid}})^2}{3} > 0.$$

Let us consider two extreme cases. First, suppose player i 's signal s_i is completely uninformative about α_i , that is, α_i and s_i are stochastically independent. The best predictor (in terms of mean squared prediction error) of α_i that is function of s_i is therefore its mean: $\mathbb{E}(\alpha_i | s_i) = \mathbb{E}(\alpha_i)$; it is constant across all states of nature because the signal conveys no information about the private

48. The σ -field generated by a signal is equal to the σ -field generated by the partition induced by the signal.

benefit parameter—it is essentially useless in predicting the value of the private benefit parameter at the true but unobservable state of nature. It follows that the variance of $\mathbb{E}(\alpha_i | s_i)$ is zero and, therefore, minimal: $\text{var}(\mathbb{E}(\alpha_i | s_i)) = \text{var}(\mathbb{E}(\alpha_i)) = 0$. Second, suppose player i 's signal s_i is completely informative about α_i , that is, $\alpha_i = f(s_i)$ for some nonconstant function $f: \mathbb{R}^3 \rightarrow \mathbb{R}_{++}$. The best predictor of α_i is therefore α_i itself: $\mathbb{E}(\alpha_i | s_i) = \mathbb{E}(f(s_i) | s_i) = f(s_i) = \alpha_i$; in other words, the player observes the value of the private benefit parameter at the true state of nature because she observes the value of her signal. It follows that the variance of $\mathbb{E}(\alpha_i | s_i)$ is equal to the variance of α_i and, therefore, maximal (see Lemma 6.2 below): $\text{var}(\mathbb{E}(\alpha_i | s_i)) = \text{var}(\alpha_i)$. Between these two extreme cases, $\text{var}(\mathbb{E}(\alpha_i | s_i))$ lies between 0 and $\text{var}(\alpha_i)$. The larger $\text{var}(\mathbb{E}(\alpha_i | s_i))$, that is, the closer $\text{var}(\mathbb{E}(\alpha_i | s_i))$ is to $\text{var}(\alpha_i)$, the more informative is the signal s_i about the private benefit parameter α_i . The quantity $\text{var}(\mathbb{E}(\alpha_i | s_i))$ is therefore a measure of the informativeness of s_i about α_i . Intuition suggests that player i will favor the second extreme case over the first when confronted with the choice between the two alternatives, that is, to be informed is better than being uninformed about the value of the private benefit parameter at the true state of nature. As will be shown below (see, in particular, Example 6.7 in Section 6.3), this is, however, in general not true.

To make the preceding observations more precise and state results, we introduce some notation and terminology. For all $i \in [I]$, let $\tilde{s}_i := (\tilde{s}_{i,\alpha}, \tilde{s}_{i,\beta}, \tilde{s}_{i,\gamma}): \Omega \rightarrow \mathbb{R}^3$ denote an alternative to player i 's signal s_i , specifically, \tilde{s}_i is a random 3-vector on the probability space $(\Omega, \mathfrak{G}, \mathbb{P})$ that has a finite support. The signal \tilde{s}_i is called *more informative* than the signal s_i if $\sigma(s_i) \neq \sigma(\tilde{s}_i)$ and $\sigma(s_i) \subset \sigma(\tilde{s}_i)$. An I -tuple of finite σ -fields on Ω is called *information structure*. Let $\sigma := (\sigma(s_1), \dots, \sigma(s_I))$ and $\tilde{\sigma} := (\sigma(\tilde{s}_1), \dots, \sigma(\tilde{s}_I))$. The information structure $\tilde{\sigma}$ is called *partially more informative for player i* than the information structure σ if $\sigma(s_i) \neq \sigma(\tilde{s}_i)$ and $\sigma(s_i) \subset \sigma(\tilde{s}_i)$ and for all $j \in [I] \setminus \{i\}$, $\sigma(s_j) = \sigma(\tilde{s}_j)$. The information structure $\tilde{\sigma}$ is called *totally more informative* than the information structure σ if for all $i \in [I]$, $\sigma(s_i) \neq \sigma(\tilde{s}_i)$ and $\sigma(s_i) \subset \sigma(\tilde{s}_i)$.⁴⁹

Definition 6.1 The *informativeness* of the signal s_i about the payoff parameter $\pi_i \in \{\alpha_i, \beta_i, \gamma_i\}$ is defined as $\text{var}(\mathbb{E}(\pi_i | s_i))$ and denoted by $\mathbb{I}(\pi_i, s_i)$.

The informativeness of a player's signal about a payoff parameter is, by definition, bounded below by zero. It is bounded above by the variance of the payoff parameter and decreasing in the distance (in the Hilbert space of square-integrable random variables on the probability space $(\Omega, \mathfrak{G}, \mathbb{P})$) between the payoff parameter and the conditional expectation of that parameter given the signal (Lemma 6.2).

Lemma 6.2 For all $\pi_i \in \{\alpha_i, \beta_i, \gamma_i\}$, $\mathbb{I}(\pi_i, s_i) = \text{var}(\pi_i) - \|\pi_i - \mathbb{E}(\pi_i | s_i)\|_2^2$.⁵⁰

The informativeness of a player's signal about a payoff parameter is a function of the joint distribution of the signal and the payoff parameter. It is zero and, therefore, minimal if the signal is completely uninformative about the payoff parameter. It is equal to the variance of the payoff parameter and, therefore, maximal if the signal is completely informative about the payoff parameter. It is increasing in the σ -field generated by the signal (Lemma 6.3).

Lemma 6.3 If $\sigma(s_i) \neq \sigma(\tilde{s}_i)$ and $\sigma(s_i) \subset \sigma(\tilde{s}_i)$, then for all $\pi_i \in \{\alpha_i, \beta_i, \gamma_i\}$, $\mathbb{I}(\pi_i, s_i) \leq \mathbb{I}(\pi_i, \tilde{s}_i)$.

49. The terminology is suggested by the notions of partial and total derivatives of a function of several variables, where a partial derivative is defined with respect to one of the variables, with the others held constant, and the total derivative allows all variables to vary.

50. Note that $\|\pi_i - \mathbb{E}(\pi_i | s_i)\|_2 = \mathbb{E}((\pi_i - \mathbb{E}(\pi_i | s_i))^2)^{1/2}$ is the distance between π_i and $\mathbb{E}(\pi_i | s_i)$.

A change in the informativeness of a player's signal about a payoff parameter is called *mean-preserving* if it does not change the expectation of the payoff parameter. A change is mean-preserving if it is caused by a change in the signal that induces a coarser or finer partition of the state space.

6.2 Dependence of ex ante expected equilibrium payoffs on informativeness of signal about payoff parameter

In the context of the Bayesian network game, ex ante expected equilibrium payoff is a natural choice for measuring the value of information. Unlike a discussion of numerical examples, a formal discussion of the value of information calls for analyzing the relation between information as measured by informativeness and ex ante expected equilibrium payoff. The findings of this analysis, paired with our understanding of the relation between information and informativeness (Lemma 6.3), form the basis for discussing the value of information.

The starting point for the analysis is a characterization of ex ante expected equilibrium payoff in the Bayesian network game \mathcal{B} with information structure σ , which is denoted by $\mathcal{B}(\sigma)$. For all $i \in [I]$, player i 's ex ante expected equilibrium payoff in $\mathcal{B}(\sigma)$ is denoted by $\mathbb{E}(u_i^*(\sigma))$ and given by

$$\begin{aligned} \mathbb{E}(u_i^*(\sigma)) &= \frac{1}{2} \mathbb{E}((\beta_i + \gamma_i)(x_i^* \circ s_i)^2) - \frac{1}{2} \sum_{m \in [I]} \bar{a}_{i,m}(D)^2 \mathbb{E}(\gamma_i(x_m^* \circ s_m)^2) \\ &\quad - \sum_{m \in [I]} \sum_{n \in [m-1]} \bar{a}_{i,m}(D) \bar{a}_{i,n}(D) \mathbb{E}(\gamma_i(x_m^* \circ s_m)(x_n^* \circ s_n)). \end{aligned} \quad (4)$$

Equation (4) shows that a player's ex ante expected equilibrium payoff is determined by the variation of her equilibrium strategy across the states of nature (see the first term on the right-hand side of equation (4)) and, provided that she is not quasi-isolated, the variations of her out-neighbors' equilibrium strategies across the states of nature (see the second term) and the interdependence of her out-neighbors' equilibrium strategies (see the third term). The variation of a player's equilibrium strategy across the states of nature is in part determined by the informativenesses of her signal about her payoff parameters. To make this dependence explicit (see Proposition 6.5 below), we introduce the following condition.

Condition 6.4 (1) For all $i \in [I]$, β_i and s_i are stochastically independent.

(2) For all $i \in [I]$ and for all $(m, n) \in (\{i\} \cup \mathcal{N}_D^+(i)) \times \mathcal{N}_D^+(i)$ with $m \neq n$, γ_i , s_m , s_n are stochastically independent.

Condition 6.4 (1) is the same as Condition 4.4 (1). Condition 6.4 (2) is stronger than Condition 4.4 (2) and has two logical consequences: First, a player's signal is completely uninformative about her own and her in-neighbors' social cost parameters. Second, the signals of a player and her out-neighbors are pairwise stochastically independent, which implies that their equilibrium strategies are uncorrelated; specifically, for all $i \in [I]$ and for all $(m, n) \in (\{i\} \cup \mathcal{N}_D^+(i)) \times \mathcal{N}_D^+(i)$ with $m \neq n$, $\text{cov}(x_m^* \circ s_m, x_n^* \circ s_n) = 0$. It does, however, not imply that they are unrelated; on the contrary, they are interdependent because of the network connecting the players; specifically, if Condition 6.4 is satisfied, then the players' ex ante expected equilibrium actions are related by

the following system of equations:

$$\mathbb{E} \left(\begin{pmatrix} x_1^* \circ s_1 \\ \vdots \\ x_I^* \circ s_I \end{pmatrix} \right) = \left(\text{diag}(\mathbb{E}(\beta)) - \text{diag}(\mathbb{E}(\gamma)) (\bar{A}(D) - E_I) \right)^{-1} \mathbb{E}(\alpha). \quad (5)$$

Proposition 6.5 *If Condition 6.4 is satisfied, then for all $i \in [I]$,*

$$\text{var}(x_i^* \circ s_i) = \frac{\mathbb{I}(\alpha_i, s_i)}{(\mathbb{E}(\beta_i + \gamma_i))^2}.$$

Proposition 6.5 shows that the variance of a player's equilibrium strategy is strictly increasing in the informativeness of her signal about her private benefit parameter if Condition 6.4 is satisfied. It implies that a player's equilibrium strategy is constant if her signal is completely uninformative about her private benefit parameter.

A player's ex ante expected equilibrium payoff is in general not monotone in the informativeness of a player's signal about a payoff parameter (see Example C.3 in Appendix C for a case where the dependence is not monotone). Condition 6.4 stipulates conditions that are sufficient for such monotone dependencies to exist.

Proposition 6.6 *Suppose Condition 6.4 is satisfied.*

(1) *For all $i \in [I]$,*

$$\begin{aligned} \mathbb{E}(u_i^*(\sigma)) &= \frac{\mathbb{I}(\alpha_i, s_i)}{2\mathbb{E}(\beta_i + \gamma_i)} - \frac{\mathbb{E}(\gamma_i)}{2} \sum_{n \in [I]} \bar{a}_{i,n}(D)^2 \frac{\mathbb{I}(\alpha_n, s_n)}{(\mathbb{E}(\beta_n + \gamma_n))^2} \\ &\quad + \frac{\mathbb{E}(\beta_i + \gamma_i)}{2} (\mathbb{E}(x_i^* \circ s_i))^2 - \frac{\mathbb{E}(\gamma_i)}{2} \left(\sum_{n \in [I]} \bar{a}_{i,n}(D) \mathbb{E}(x_n^* \circ s_n) \right)^2. \end{aligned} \quad (6)$$

(2) *A mean-preserving decrease (respectively, increase) in the informativeness of a player's signal about her private benefit parameter strictly decreases (respectively, increases) her ex ante expected equilibrium payoff and strictly increases (respectively, decreases) an in-neighbor's ex ante expected equilibrium payoff if the in-neighbor is not quasi-isolated.*

Result 1 of Proposition 6.6 provides details on the relation between a player's ex ante expected equilibrium payoff and the informativeness of her signal about her private benefit parameter and those of her in-neighbors about their private benefit parameters. When interpreting this relation, it is important to bear in mind that the informativeness of a signal about a payoff parameter is not per se a parameter but a function of the joint distribution of the signal and the payoff parameter, and a change in the joint distribution may affect determinants of ex ante expected equilibrium payoff other than the informativeness, for example, ex ante expected equilibrium action. It is, therefore, in general not always appropriate to assume a ceteris paribus change in the informativeness of a player's signal about a payoff parameter to study its effects on her and other players' ex ante expected equilibrium payoffs.

The monotonicities asserted by Result 2 of Proposition 6.6 do not necessarily hold true if the change in the informativeness is not mean-preserving. The reason for this is that a change in

the informativeness of a player's signal about her private benefit parameter changes her and possibly other players' ex ante expected equilibrium actions if the change is not mean-preserving (see the system of equations (5)). The changes in ex ante expected equilibrium actions may in turn decrease or increase a player's ex ante expected equilibrium payoff or leave it unchanged (see the third and the fourth terms on the right-hand side of equation (6)).

6.3 The value of private information

To study the value of private information, we increase—*ceteris paribus*—the informational content of a player's signal and analyze its effect on her and other players' ex ante expected equilibrium payoffs. More specifically, we consider two Bayesian network games $\mathcal{B}(\sigma)$ and $\mathcal{B}(\tilde{\sigma})$ that differ only in their information structures σ and $\tilde{\sigma}$, one of which is partially more informative for a single player than the other.

More information is in general not beneficial for a player, that is, it may have a positive value, no value, or a negative value in terms of ex ante expected equilibrium payoff. While the topology of the network may account for a negative value of more information, as illustrated by the example in Section A.5.1 (see Appendix A), it is not necessarily the sole cause. Even a quasi-isolated player, whose equilibrium strategy and ex post expected equilibrium payoff are given by the formulae in (3), may prefer less information to more information, as illustrated by Example 6.7.

Example 6.7 Suppose player i is quasi-isolated, her signal in $\mathcal{B}(\sigma)$ is completely uninformative about α_i and β_i , and her signal in $\mathcal{B}(\tilde{\sigma})$ is completely informative about α_i and β_i . It follows that

$$\mathbb{E}(u_i^*(\tilde{\sigma})) - \mathbb{E}(u_i^*(\sigma)) = \frac{1}{2} \text{cov}\left(\alpha_i^2, \frac{1}{\beta_i}\right), \quad (7)$$

which is negative (respectively, positive) if α_i and β_i are positively (respectively, negatively) correlated. \diamond

General statements about the value of more information are, nonetheless, possible if certain conditions are met by the structure of the Bayesian network game, including Condition 6.4. The results are stated in Corollary 6.8.

Corollary 6.8 *Suppose Condition 6.4 is satisfied in $\mathcal{B}(\sigma)$ and $\mathcal{B}(\tilde{\sigma})$ and $\tilde{\sigma}$ is partially more informative for player i than σ .*

(1) $\mathbb{E}(u_i^*(\tilde{\sigma})) \geq \mathbb{E}(u_i^*(\sigma))$.

(2) *If player n is an in-neighbor in D of player i and not quasi-isolated, then $\mathbb{E}(u_n^*(\tilde{\sigma})) \leq \mathbb{E}(u_n^*(\sigma))$.*

Results 1 and 2 hold true with strict inequalities if $\mathbb{I}(\alpha_i, \tilde{s}_i) > \mathbb{I}(\alpha_i, s_i)$.

Under the conditions of Corollary 6.8, private information has a positive value to a player and imposes a negative externality on her in-neighbors if it increases the informativeness of her signal about her private benefit parameter.

6.4 The value of public information

To study the value of public information, we supplement the structure of the Bayesian network game \mathcal{B} with a public signal. The resulting network game is called an *extension* of \mathcal{B} , wherein all players observe the values of their signals and the value of the public signal. The public signal is a random variable $s_p: \Omega \rightarrow \mathbb{R}$ on the probability space $(\Omega, \mathfrak{S}, \mathbb{P})$ that has a finite support; it can be thought of as constituting the fourth component of a player's signal, so no extra notation is required.

Let $\bar{\mathcal{B}}(\sigma)$ and $\bar{\mathcal{B}}(\tilde{\sigma})$ denote two extensions of \mathcal{B} whose structures are the same except for their information structures σ and $\tilde{\sigma}$. An extension of \mathcal{B} is strategically equivalent to \mathcal{B} if the public signal in the extension is constant. The information structure $\tilde{\sigma}$ is totally more informative than the information structure σ if and only if the public signal in $\bar{\mathcal{B}}(\tilde{\sigma})$ is more informative than the public signal in $\bar{\mathcal{B}}(\sigma)$. Suppose $\tilde{\sigma}$ is totally more informative than σ . The difference between player i 's ex ante expected equilibrium payoffs in $\bar{\mathcal{B}}(\tilde{\sigma})$ and $\bar{\mathcal{B}}(\sigma)$, $\mathbb{E}(u_i^*(\tilde{\sigma})) - \mathbb{E}(u_i^*(\sigma))$, represents the value to player i of the informational content of the public signal in $\bar{\mathcal{B}}(\tilde{\sigma})$ —in other words, the value of public information—if the public signal in $\bar{\mathcal{B}}(\sigma)$ is constant, otherwise it represents the value of the difference in informational contents of the public signals in $\bar{\mathcal{B}}(\tilde{\sigma})$ and $\bar{\mathcal{B}}(\sigma)$ —in other words, the value of more public information.

Public information or more public information is in general not beneficial for a player, that is, depending on the structure of the Bayesian network game, it may have a positive value, no value, or a negative value; see the example in Section A.5.2 in the Appendix for an illustration.

7 Concluding remarks

We develop a network model in which players bear a cost from deviating from the social norm of their peers and know with uncertainty different parameters of their utility function, such as the marginal costs and benefits of committing crime. We study how the social environment, and the idiosyncratic characteristics of the players, affect their equilibrium actions and welfare. One interesting result is to show that denser networks do not necessarily increase agents' actions and welfare. We also show that an upward shift in the idiosyncratic component of a given player's taste for conformity leads to an increase in this player's action if her action is greater than that of her social norm. We also find that, under some conditions, it is optimal for the planner to affect the payoffs of selected individuals rather than all agents in the network. We finally show that having more information is not always beneficial to agents and can, in fact, reduce their welfare.

We illustrate all our results in the context of criminal networks in which offenders do not know with certitude the probability of being caught and the severity of the judges if arrested and do not want to be different from their peers in terms of criminal activities. It should be clear that we could use our framework to study other outcomes as long as agents do not know with certainty the private benefits and costs of their actions, are embedded in a network and pay a cost from deviating from the social of their peers.

Consider, for example, tax evasion. Individuals who want to tax evade have to decide how much income they have to declare but do not know the probability of being audited. They are also strongly influenced by their peers (see, e.g., Alm, Bloomquist, and McKee (2017), Galbiati and Zanella (2012) and Fortin, Lacroix, and Villeval (2007)). Environment practices such as recycling can also be another nice application of our model. We know that social norms matter and that people do know the exact benefit of having pro-environmental behavior (see e.g. Farrow,

Grolleau, and Ibanez (2017)).

More generally, we believe that our model contributes to the theory of games on networks with incomplete information and sheds some light on how different policies can be implemented to increase welfare in different activities. These issues are complex and we hope that more research will be undertaken in the future.

References

- Akerlof, G. A. 1997. "Social Distance and Social Decisions." *Econometrica* 65 (5): 1005–1027.
- Alm, J., K. Bloomquist, and M. McKee. 2017. "When you Know your Neighbour Pays Taxes: Information, Peer Effects, and Tax Compliance." *Fiscal Studies* 38:587–613.
- Apel, R. 2013. "Sanctions, Perceptions, and Crime: Implications for Criminal Deterrence." *Journal of Quantitative Criminology* 29 (1): 67–101.
- Ballester, C., A. Calvó-Armengol, and Y. Zenou. 2006. "Who's Who in Networks. Wanted: The Key Player." *Econometrica* 74 (5): 1403–1417.
- . 2010. "Delinquent Networks." *Journal of the European Economic Association* 8 (1): 34–61.
- Bang-Jensen, J., and G. Z. Gutin. 2009. *Digraphs: Theory, Algorithms and Applications*. 2nd ed. Springer Monographs in Mathematics. London: Springer-Verlag.
- Baumann, F., and T. Friehe. 2013. "Cheap Talk About the Detection Probability." *International Game Theory Review* 15 (1): 1–16.
- Bebchuk, L., and L. Kaplow. 1992. "Optimal Sanctions When Individuals Are Imperfectly Informed about the Probability of Apprehension." *The Journal of Legal Studies* 21 (2): 365–370.
- Becker, G. 1968. "Crime and punishment: an economic approach." *Journal of Political Economy* 76 (2): 169–217.
- Bergemann, D., T. Heumann, and S. Morris. 2015. "Networks, Information and Volatility." Unpublished manuscript, Yale University.
- Berman, A., and R. J. Plemmons. 1994. *Nonnegative Matrices in the Mathematical Sciences*. Classics in Applied Mathematics. Philadelphia: Society for Industrial and Applied Mathematics.
- Bhuller, M., G. B. Dahl, K. V. Løken, and M. Mogstad. 2018. "Incarceration Spillovers in Criminal and Family Networks." Unpublished manuscript, University of Chicago.
- Blanes i Vidal, J., and G. Mastrobuoni. 2018. "Police Patrols and Crime." IZA DP No. 11393.
- Blume, L. E., W. A. Brock, S. N. Durlauf, and R. Jayaraman. 2015. "Linear Social Interactions Models." *Journal of Political Economy* 123 (2): 444–496.
- Boucher, V. 2016. "Conformism and Self-Selection in Social Networks." *Journal of Public Economics* 136:30–44.

- Bramoullé, Y., and R. Kranton. 2016. "Games Played on Networks." In *The Oxford Handbook of the Economics of Networks*, edited by Y. Bramoullé, B. W. Rogers, and A. Galeotti. Oxford University Press.
- Bramoullé, Y., R. Kranton, and M. D'Amours. 2014. "Strategic Interaction and Networks." *American Economic Review* 104 (3): 898–930.
- Calvó-Armengol, A., and J. de Martí. 2007. "Communication Networks: Knowledge and Decisions." *American Economic Review Papers and Proceedings* 97 (2): 86–91.
- . 2009. "Information Gathering in Organizations: Equilibrium, Welfare and Optimal Network Structure." *Journal of the European Economic Association* 7:116–161.
- Calvó-Armengol, A., J. de Martí, and A. Prat. 2015. "Communication and Influence." *Theoretical Economics* 10:649–690.
- Calvó-Armengol, A., and Y. Zenou. 2004. "Social Networks and Crime Decisions. The Role of Social Structure in Facilitating Delinquent Behavior." *International Economic Review* 45:939–958.
- Chalfin, A., and J. McCrary. 2017. "Criminal Deterrence: A Review of the Literature." *Journal of Economic Literature* 55 (1): 5–48.
- Cortés, D., G. Friebel, and D. Maldonado. 2019. "Crime and Education in a Model of Information Transmission." *Annals of Public and Cooperative Economics*, forthcoming.
- Damm, A. P., and C. Dustmann. 2014. "Does Growing Up in a High Crime Neighborhood Affect Youth Criminal Behavior?" *American Economic Review* 104 (6): 1806–1832.
- Farrow, K., G. Grolleau, and L. Ibanez. 2017. "Social Norms and Proenvironmental Behavior: A Review of the Evidence." *Ecological Economics* 140:1–13.
- Fortin, B., G. Lacroix, and M.-C. Villeval. 2007. "Tax evasion and social interactions." *Journal of Public Economics* 91 (11): 2089–2112.
- Frobenius, F. G. 1912. "Über Matrizen aus nicht negativen Elementen." *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften zu Berlin*: 456–477.
- Frommer, A. 1990. *Lösung linearer Gleichungssysteme auf Parallelrechnern*. Braunschweig: Vieweg Verlag.
- Galbiati, R., and G. Zanella. 2012. "The Tax Evasion Social Multiplier: Evidence from Italy." *Journal of Public Economics* 96:485–494.
- Galeotti, A., C. Ghiglino, and F. Squintani. 2013. "Strategic Information Transmission Networks." *Journal of Economic Theory* 148 (5): 1751–1769.
- Galeotti, A., S. Goyal, M. O. Jackson, F. Vega-Redondo, and L. Yariv. 2010. "Network Games." *The Review of Economic Studies* 77 (1): 218–244.
- Gardner, M., and L. Steinberg. 2005. "Peer Influence on Risk Taking, Risk Preference, and Risky Decision Making in Adolescence and Adulthood: An Experimental Study." *Developmental Psychology* 41 (4): 625–635.

- Glaeser, E. L., B. Sacerdote, and J. A. Scheinkman. 1996. "Crime and Social Interactions." *Quarterly Journal of Economics* 111 (2): 507–548.
- Golub, B., and S. Morris. 2017. "Higher-Order Expectations." Unpublished manuscript, Harvard University.
- . 2018. "Expectations, Networks, and Conventions." Unpublished manuscript, Harvard University.
- Hagenbach, J., and F. Koessler. 2010. "Strategic Communication Networks." *The Review of Economic Studies* 77:1072–1099.
- Hervés-Beloso, C., and P. Klinger Monteiro. 2013. "Information and σ -algebras." *Economic Theory* 54 (2): 405–418.
- Hjalmarsson, R. 2008. "Crime and Expected Punishment: Changes in Perceptions at the Age of Criminal Majority." *American Law and Economics Review* 11 (1): 209–248.
- Jackson, M. O. 2008. *Social and Economics Networks*. Princeton: Princeton University Press.
- Jackson, M. O., and Y. Zenou. 2015. "Games on Networks." Chap. 3 in *Handbook of Game Theory with Economic Applications*, edited by H. P. Young and S. Zamir, 4:95–163. Elsevier.
- Klenke, A. 2014. *Probability Theory: A Comprehensive Course*. 2nd ed. Universitext. London: Springer-Verlag.
- Kling, J. R., J. Ludwig, and L. F. Katz. 2005. "Neighborhood Effects on Crime for Female and Male Youth: Evidence from a Randomized Housing Voucher Experiment." *Quarterly Journal of Economics* 120 (1): 87–130.
- Leister, C. M. 2017. "Information Acquisition and Welfare in Network Games." Unpublished manuscript, Monash University.
- Leister, C. M., Y. Zenou, and J. Zhou. 2019. "Coordination on Networks." Unpublished manuscript, Monash University.
- Lindquist, M. J., and Y. Zenou. 2014. "Key Players in Co-Offending Networks." CEPR Discussion paper No. 9889.
- . 2019. "Crime and networks: 10 policy lessons." *Oxford Review of Economic Policy* forthcoming.
- Liu, X., E. Patacchini, Y. Zenou, and L.-f. Lee. 2012. "Criminal networks: Who is the key player?" CEPR Discussion paper No. 8772.
- Lochner, L. 2007. "Individual Perceptions of the Criminal Justice System." *American Economic Review* 97 (1): 444–460.
- Marschak, J., and R. Radner. 1972. *Economic Theory of Teams*, vol. Monograph 22. Cowles Foundation for Research in Economics at Yale University. New Haven and London: Yale University Press.
- Martí, J. de, and Y. Zenou. 2015. "Network Games with Incomplete Information." *Journal of Mathematical Economics* 61:221–240.

- McDonald, J., M. Neumann, H. Schneider, and M. Tsatsomeros. 1995. "Inverse M-matrix inequalities and generalized ultrametric matrices." *Linear Algebra and its Applications* 220:321–341.
- Meyer, C. D. 2000. *Matrix Analysis and Applied Linear Algebra*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics.
- Myatt, D., and C. Wallace. 2019. "Information Acquisition and Use by Networked Players." *Journal of Economic Theory* 182:360–401.
- Parker, J., and H. G. Grasmick. 1979. "Linking actual and perceived certainty of punishment: An exploratory study of an untested proposition in deterrence theory." *Criminology* 17 (3): 366–379.
- Patacchini, E., and Y. Zenou. 2012. "Juvenile delinquency and conformism." *Journal of Law, Economics, and Organization* 28 (1): 1–31.
- Perron, O. 1907. "Zur Theorie der Matrices." *Mathematische Annalen* 64 (2): 248–263.
- Polinsky, A. M., and S. Shavell. 2000. "The Economic Theory of Public Enforcement of Law." *Journal of Economic Literature* 38 (1): 45–76.
- Radner, R. 1968. "Competitive Equilibrium Under Uncertainty." *Econometrica* 36 (1): 31–58.
- Sah, R. K. 1991. "Social osmosis and patterns of crime." *Journal of Political Economy* 99 (6): 1272–1295.
- Schrag, J., and S. Scotchmer. 1997. "The self-reinforcing nature of crime." *International Review of Law and Economics* 17 (3): 325–335.
- Sommer, M. 2017. *Essays on Games on Networks*. Zurich: University of Zurich.
- Ushchev, P., and Y. Zenou. 2019. "Social Norms in Networks." *Journal of Economic Theory*, forthcoming.
- Varga, R. S. 2000. *Matrix Iterative Analysis*. 2nd ed. Vol. 27. Springer Series in Computational Mathematics. Berlin and Heidelberg, Germany: Springer-Verlag.
- Zenou, Y. 2016. "Key Players." In *The Oxford Handbook of the Economics of Networks*, edited by Y. Bramoullé, B. W. Rogers, and A. Galeotti, 244–276. Oxford University Press.