



Universiteit
Leiden
The Netherlands

The desperation threshold: a model to explain decisions in poverty

Du Buisson de Courson, B.

Citation

Du Buisson de Courson, B. (2026, April 7). *The desperation threshold: a model to explain decisions in poverty*. Retrieved from <https://hdl.handle.net/1887/4301048>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4301048>

Note: To cite this publication please use the final published version (if applicable).

2 Why is violence high and persistent in deprived communities? A formal model

This chapter is based on: de Courson, B., Frankenhuis, W. E., Nettle, D., & Van Gelder, J. L. (2023). Why is violence high and persistent in deprived communities? A formal model. *Proceedings of the Royal Society B*, 290(1993), 20222095.

2.1 Abstract

There is massive variation in rates of violence across time and space. These rates are positively associated with economic deprivation and inequality. They also tend to display a degree of local persistence, or ‘enduring neighbourhood effects’. Here, we identify a single mechanism that can produce all three observations. We formalise it in a mathematical model, which specifies how individual-level processes generate the population-level patterns. Our model assumes that agents try to keep their level of resources above a ‘desperation threshold’, to reflect the intuitive notion that one of people’s priorities is to always meet their basic needs. As shown in previous work, being below the threshold makes risky actions, such as property crime, beneficial. We simulate populations with heterogeneous levels of resources. When deprivation or inequality is high, there are more desperate individuals, hence a higher risk of exploitation. It then becomes advantageous to use violence, to send a “toughness signal” to exploiters. For intermediate levels of poverty, the system is bistable and we observe hysteresis: populations can be violent because they were deprived or unequal in the past, even after conditions improve. We discuss implications of our findings for policy and interventions aimed at reducing violence.

2.2 Introduction

There is massive variation in neighbourhood levels of interpersonal violence across time (Eisner, 2003) and space (Glaeser et al., 1995). For instance, homicide rates varied more than 100-fold between Chicago neighbourhoods in the 1988-93 period (Wilson & Daly, 1997). This variation is partially explained by macro-level factors (Pratt & Cullen, 2005), such as economic conditions, poverty and inequality in particular (Daly, 2017; Fajnzylber et al., 2002; Hsieh & Pugh, 1993; Kelly, 2000; Pratt & Cullen, 2005; M. Wilson & Daly, 1997), with extremely disadvantaged neighbourhoods having an unusually high level of violent crime (Krivo & Peterson, 1996). Violence is also a locally persistent phenomenon, subject to enduring neighbourhood effects (Sampson, 2012). That is, violent places tend to remain violent to some degree, despite continuous population flux and even when economic conditions improve (Sampson, 2012; Shaw & McKay, 1942; Weisburd et al., 2004).

Different bodies of research have offered explanations of these observations at different scales. At the population level, several criminological theories view poverty as interfering with the normal functioning of a community, creating ‘strain’ (Agnew, 1992; Messner & Rosenfeld, 1994). Other criminological theories propose that poverty renders neighbourhoods deficient in social organisation (Shaw & McKay, 1942) or social efficacy (Sampson, 2012), promoting violence and crime. A separate body of work has focused on the role of inequality, arguing that inequality creates a fiercer competition for symbolic and material resources, resulting in higher violence (Daly, 2017; R. Wilkinson, 2004; R. G. Wilkinson et al., 1998; R. Wilkinson & Pickett, 2010).

At the individual and psychological levels, several authors have pointed to the role of time preferences (Brezina et al., 2009; M. Wilson & Daly, 1997). Poverty (Bolte et al., 2010; Lantz et al., 1998) and inequality (M. Wilson & Daly, 1997) are related to poorer future prospects, including higher mortality and morbidity rates. This can result in a sense of futurelessness, which in turn leads to steep future discounting and choosing actions that can lead to immediate payoffs, such as crime (Brezina et al., 2009; Daly & Wilson, 2001). Yet,

other authors report that violence is a signal that serves to communicate a toughness reputation and avoid being victimised. This idea has been proposed independently in a variety of fields: ethnography (Anderson, 2000; Brezina et al., 2016), sociology (Gambetta, 2009), cultural psychology (D. Cohen & Nisbett, 2016) and evolutionary psychology (Fessler et al., 2014). On this view, violence also has long-term benefits, and does not necessarily qualify as a short-term strategy. Therefore, time preferences alone can not explain the social gradient of violence.

In this article, we show that a single mechanism is able to generate all three key observations: violence is higher in deprived or unequal populations, varies considerably between populations, and can persist in a community despite economic improvement. Whereas the theories mentioned above are stated verbally, we articulate our explanation in a formal model. This serves two purposes. The first is to prove that the mechanism we propose is indeed able to reproduce the empirical observations at a qualitative level (i.e. the model's generative sufficiency (J. M. Epstein, 2012). If it is, our model can be considered as a valid candidate explanation (J. M. Epstein, 2012, 2012). Second, formalisation eliminates the ambiguity inherent in natural language (Frankenhuis et al., 2023) and compels the provision of a fully explicit mechanism (Smaldino, 2017; Smaldino, 2020). In particular, the process of formalisation forces the specification of how interactions at the individual level produce group-level outcomes, which in turn shape individual behaviour. Thus, our approach aligns with the key aim of criminology and the social sciences generally to integrate micro- and macro-level processes (Box-Steffensmeier et al., 2022; Matsueda, 2017). This model uses ideas from complexity science. In that field, several models have studied crime (for a review, see D'Orsogna & Perc, 2015) – yet without engaging with the role of material circumstances.

That violence is a social problem (Heeks et al., 2018) does not imply that it reflects a dysfunction at the individual level. Here, we propose that violence is a 'contextually appropriate response', meaning it can be understood as a response to the costs and benefits associated with living in a particular context – as opposed to, for instance, a psychopathology or failure of willpower (Pepper & Nettle, 2017). In our model, individuals make decisions based on their level of resources and other individuals' behaviours. This game-theoretical feature creates the possibility that optimal strategies at the individual level produce suboptimal outcomes for the population, such as high rates of community violence, analogously to a 'tragedy of the commons' (Hardin, 1968).

Analysing violence as a contextually appropriate response requires first the specification of its possible costs and benefits for an individual. Violence is commonly assumed – for instance, in the classic Hawk-Dove model (Maynard Smith, 1982) – to allow an agent to take a resource by force, while facing a physical risk. In addition to these assumptions, we incorporate the idea that violence has reputational consequences as a 'toughness signal' (Anderson, 2000; Brezina et al., 2016; D. Cohen & Nisbett, 2016; Fessler et al., 2014; Gambetta, 2009), reducing the probability of being exploited. We thus focus on interpersonal violence involving physical harm to others for instrumental or reputational motives, rather than other forms of violence, like self-harm, child abuse, intimate partner violence, or warfare. Our question then becomes "why would it be more appropriate to send such signals in deprived or unequal neighbourhoods?". These neighbourhoods can be characterised as having a larger number of 'desperate individuals', without enough resources to meet their basic needs. They can be compared to drowning individuals, who would do anything to try to get their head out of the water, including dragging others down. Despite the high potential costs (de Courson & Nettle, 2021), exploiting others can be the most direct way to get resources quickly and jump back "above water". We argue that the risk of being exploited by desperate individuals, in turn, triggers an incentive to send toughness signals – i.e., violent displays – among non-desperate individuals, to protect themselves from being targeted.

To formalise this intuition, we make two main assumptions. To represent the ‘signal effect’ of violence, we assume that being violent makes one less likely to be selected as a target of property crime. To represent financial desperation, we assume that agents are defined by a dynamic level of resources and have a ‘desperation threshold’, below which it is harmful to fall. In other words, agents are trying to always meet their basic needs and keep their head above water. The assumption of a threshold is a theoretically innovative idea in the social sciences (de Courson & Nettle, 2021) that we believe to be reasonable. It is inspired by optimal foraging theory in ethology (Stephens, 1981), where models commonly include a ‘starvation threshold’ below which fitness rapidly declines. In humans, there are ethnographic descriptions of such thresholds (Du Bois, 2010; Nettle, 2015; Scott, 1977), involving both physiological needs (e.g., hunger) and social needs (e.g., being respected), thus more generally the ability to meet basic needs. In the Philadelphia community he studied, Du Bois described the poorest individuals as a “submerged tenth”, who are more likely to engage in dangerous actions (Du Bois, 2010). Scott (1977) later observed among South Asian farmers a “subsistence crisis level—perhaps a ‘danger zone’ rather than a ‘level’ would be more accurate [...] a threshold below which the qualitative deterioration in subsistence, security and status is massive and painful” (p. 17). Experimental games have found that humans adjust their levels of risk taking in response to such thresholds (Mishra & Lalumière, 2010; Pietras & Hackenberg, 2001), including by stealing resources from other participants (Radkani et al., 2023). Here, we assess the explanatory power of the desperation threshold for the socio-economic gradient of violence.

2.3 Model

The model is more thoroughly described in the Supplementary Materials. The Python code can be found here: <https://github.com/regicid/model>

2.3.1 Structure of the model

Our model combines an individual-level optimal decision model and a population structure. The individual-level component is a state-dependent optimisation algorithm, implemented by stochastic dynamic programming (Houston et al., 1999; Mangel & Clark, 1988). Agents are defined by a dynamic ‘state variable’ that represents their level of resources. It is affected by the agent’s actions and by random fluctuations, following an AR(1) process with autocorrelation r . The model assumes a ‘desperation threshold’, a value below which agents are heavily penalised. They have access to several strategies, defined by probabilistic consequences on their level of resources, conditional on other agents’ strategies. Agents choose the strategy associated with the highest ‘fitness’, a maximand criterion that represents the agents’ goal. We allow the chosen strategy to depend on (i) the agent’s level of resources and (ii) the frequencies of strategies in the population. In other words, agents pick the strategy which is optimal for their current level of resources and social environment.

The individual-level model identifies the optimal action an individual should choose for any given distribution of strategies in the surrounding population. It is not sufficient for revealing how that distribution will evolve. To address this population-level question, we simulate large populations of interacting agents with different levels of resources. By varying the initial distribution of those resource levels – in particular their mean and variance, representing economic affluence and inequality, respectively – we can test how the economic context affects the level of violence. We run the simulations until the system reaches a stable equilibrium. We also vary the initial distributions of strategies to examine the possibility of hysteresis, the dependence of the

outcome on the initial conditions of the system.

2.3.2 Individual strategies and fitness

At each time step, agents can choose between three strategies: ‘exploitation’, ‘violence’ and ‘submission’. Exploitation represents property crime. When exploiting, an agent tries to take resources from another agent. This strategy entails two potential costs. First, the exploiter might be caught and sanctioned with a probability γ , costing π units of resources (see Table 2.1 for a summary of the parameters). This reflects exogenous social control, such as policing. We set the probabilities and magnitude of these costs (γ , and π) such that the expected payoff of stealing is always negative, ensuring exploitation stays on average a bad decision. Second, the target agent may react violently: agents either fight (‘violence’ strategy) or not do so (‘submission’). ‘Exploiters’ are also assumed to react violently to exploitation, but henceforth we use ‘violent’ to describe an agent who plays the violent strategy but does not exploit. The fight’s winner, selected by a coin toss, obtains or keeps the disputed resources. The loser pays a proportionate fitness cost λ , described in the next subsection. We also assume that violent agents sometimes enter unnecessary fights: with probability m , they attack a non-exploiter agent by mistake. This triggers a fight if the other individual is also violent.

Table 2.1. Notation summary

Symbol	Meaning	Typical value or range
N	Population size	105
μ	Mean resource level	[5, 25]
σ	Variance of resource levels	[4, 10]
r	Resource levels autocorrelation	.99
n	Number of possible targets	[1, 50]
β	Exploitation resource stake	10
π	Resource cost of punishment	20
γ	Probability of punishment	1/3
m	Probability of violent mistake	[.01, .3]
ω	Fitness cost of being below threshold	[.01, .3]
λ	Lost fight fitness cost	[.01, .3]

The violent strategy also sends a ‘toughness signal’: it is observable and confers a ‘toughness reputation’, which reduces the probability of being the target of exploitation at this round. We assume that choosing the violent strategy suffices to confer this reputation, even for an agent who has never fought. This is a simplifying assumption, that also guarantees the coherence of the model (see Supplementary Materials). Exploiters choose their target out of a set of n randomly drawn agents, among which they prefer submissive targets over violent ones. Agents playing ‘violent’ at this round therefore have (if $n > 1$) a non-zero but lower probability

of being victimised than agents playing ‘submissive’ (see Supplementary Materials for details). The parameter n controls how much lower, and thus represents the toughness signal’s efficiency: it has no effect if $n=1$, whereas if $n \rightarrow \infty$, then violent agents are never victimised as long as there is at least one submissive agent in the population. The ‘violent’ strategy is designed to represent forms of violence involving either a material or a reputational stake, our model thus focuses on these forms.

Agents choose the strategy that yields the highest expected ‘fitness’, defined as the resource level attained after T periods reduced by (i) a fraction ω for every period spent below the desperation threshold and (ii) a fraction λ for every lost fight, representing the risk of an injury when fighting. In other words, agents try to maintain their head above water financially while fighting as rarely as possible. Since we use a Stochastic Dynamic Programming algorithm, the decision takes into account not only the possibility of being below the desperation threshold immediately, but also at any later point in the future.

2.3.3 Population simulations

We simulate populations of $N = 10^5$ agents whose level of resources are drawn from a Gaussian distribution. At every time step, we let 10% of the population, randomly chosen, update their strategies. This way, the distribution of strategies in the population can evolve smoothly to an equilibrium rather than oscillating. When updating their strategies, agents take into account their level of resource and the current frequency of ‘exploitation’ and ‘violence’ in the population. For simplicity, we assume agents have perfect knowledge of others’ strategies. We iterate this process enough times for the distribution of strategies to stabilise.

We then compare the outcomes of the model for different values of the mean and the variance of the distribution of resources in the population, representing economic affluence and inequality, respectively. To investigate the possibility of hysteresis, we test if the outcome depends on the initial proportion of violent agents. In the Supplementary Materials, we explore how the model is affected by varying the other parameters.

2.4 Results

2.4.1 Individual decisions




Figure 2.1: Optimal strategies depending on resource levels (x-axis), prevalence of violence (y-axis) and prevalence of stealing (panel). The central tick on the x-axis (o) represents the desperation threshold. Agents tend to exploit below the threshold and to be violent either when exploitation is frequent, when violence is very rare or when violence is frequent.

Figure 2.1 represents the optimal decisions depending on resource levels and the prevalence of exploitation and violence in the population. It is generally optimal to exploit when one is below the desperation threshold, confirming the previous finding that being under water makes risk-taking contextually appropriate (de Courson & Nettle, 2021). The area where agents exploit is roughly a square. This decision is thus virtually independent of the level of violence, even though exploitation is greatly disincentivized by the presence of violent agents. In a fully violent population, an exploiter will always have to fight, which halves its probability of success and reduces its “fitness” in case of loss.

Above the threshold, agents should be violent in several cases. First, when exploitation is frequent (in figure 2.1, compare the right panel to the left). This is unsurprising: violence lowers the risk of being exploited, so the higher this risk is, the more agents should protect themselves. Second, when violence is very rare (in figure 2.1, at the bottom of the y axis). A lone violent agent will never be exploited (as the exploiters will always be able to choose a submissive target instead) and never fight, as it never meets other violent agents. Therefore, being violent in a fully non-violent environment has benefits but no costs, and is favoured.

Agents should also choose violence when violence is very frequent (in figure 2.1, at the top of the y-axis). The more violence there is, the more exploitation is concentrated on the rare submissive agents. This incentivizes them to be violent, and thus violence begets violence. Finally, agents are prone to violence when close to the threshold. This is due to risk preferences: being violent reduces the risk of being exploited. Just as agents are risk-prone below the threshold because they have ‘nothing to lose’, they are risk-averse just above the threshold as they have ‘too much to lose’, like a person on the edge of a cliff. Intuitively, close-to-the-edge individuals choose to risk their health in an attempt to hang on to their resources.

2.4.2 Population simulations

At the population level, outcomes depend chiefly on the proportion of agents below the desperation threshold, the ‘desperation rate’. This rate determines the number of exploiters, which in turn determines the number of violent agents (see figure 2.1). Both poverty (low mean level of resources μ) and inequality (high variance σ) increase the desperation rate: the lower or the broader the distribution, the larger the left tail consisting of desperate agents (figure 2.2).

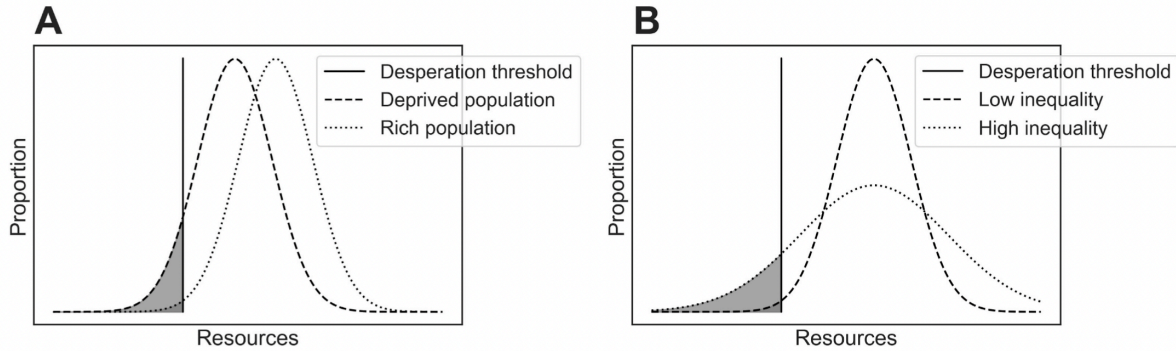


Figure 2.2: Effect of poverty (A) and inequality (B) on desperation rate. Either reducing the mean of the resource distribution or increasing its variance leads to a larger tail of individuals whose resources are below the desperation threshold.

The higher the proportion of desperate agents, the higher the prevalence of violence at equilibrium (figure 2.3). In our model, therefore, both poverty (μ) and inequality (σ) increase violence. However, the relationship of the desperation rate to the equilibrium frequency of violence depends on how we initialise the strategies. If we begin with no violent agents in the population, we reach the equilibrium frequencies shown with the crosses on figure 2.3. We observe an inverse s-shaped function. For low values of desperation rate, we have a concave relation: the violent strategy gets costlier as it spreads, dampening the increase. For high values, we observe a convex relation revealing a positive feedback effect: exploitation is deflected onto the submissive agents, and violence begets violence. In the Supplementary Materials, we show that this result is qualitatively robust to changes in the exogenous parameters values.

Here, it must be noted that inequality only plays a role by increasing the amount of desperate agents. In other words, inequality increases violence through absolute poverty, not relative poverty. Concretely, enriching the rich without impoverishing the poor would not increase violence. It must however be noted that we assumed the desperation threshold to be fixed and independent on the average level of resource, which, by construction, precludes the role of relative deprivation. We explore this limitation in the Discussion.

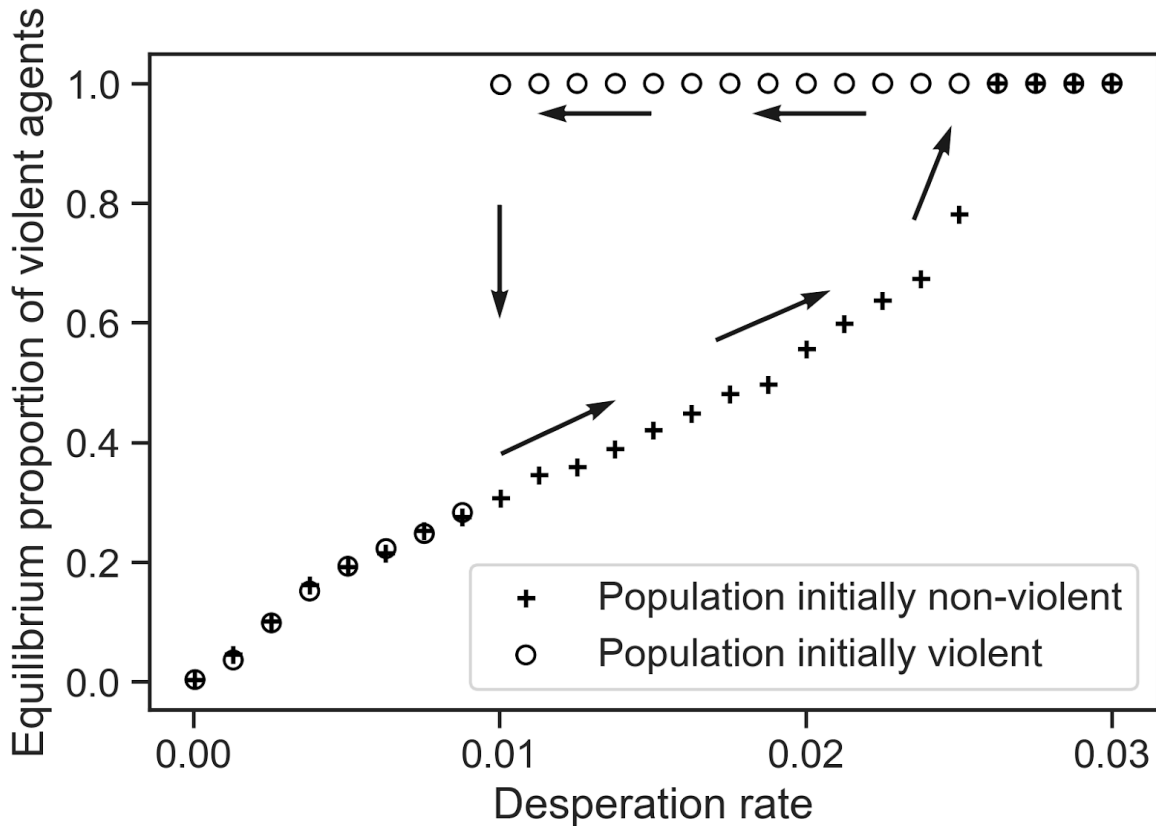


Figure 2.3: Proportion of agents playing ‘violent’ or ‘exploiting’ at equilibrium against the desperation rate. The crosses show the data for populations initialised with no violent strategies, and the circles populations initialised with all violent agents. We cut the x-axis when both curves reach 1 (beyond .03, crosses and circles stay at 1). The arrows illustrate the hysteresis loop: if a population’s desperation rate increases from 0 to .04, it would follow the lower branch, but if it decreases from .04 to 0, it would follow the upper branch.

2.4.3 Multiple equilibria and hysteresis

If we initialise the simulation with only violent agents, the population moves to a completely violent configuration in a large range of the desperation rate range (figure 2.3, circles). Here, the system has two distinct stable equilibria. Which of these is reached depends on where the system comes from. We observe a hysteresis loop: if the desperation rate (here exogenous) increases from 0 to .03, the system follows the lower branch, but if it decreases from .03 to 0, it follows the upper branch and remains in a fully violent configuration for a long time. Thus, two equally deprived or unequal communities can have vastly different levels of violence for historical reasons, with high desperation in the past producing a persisting high violence. We can visualise this bistability using a vector field (figure 2.4A). The upper equilibria have small basins of attraction, and will therefore not be reached unless a very large share of the population is violent. However, figure 2.4B shows that

the incentive to be violent soars as the prevalence of violence goes to 1, as rare non-violent agents concentrate exploitation costs. Thus, even though a few non-violent agents would suffice to reach the basin of attraction of the lower equilibrium, the model predicts these equilibria to be stable and robust to small changes in the parameters. In the Supplementary Materials, we show that the hysteresis effect holds if the parameters r and n are high enough, that is, if toughness signals are efficient enough and the experience of desperation sufficiently persistent.

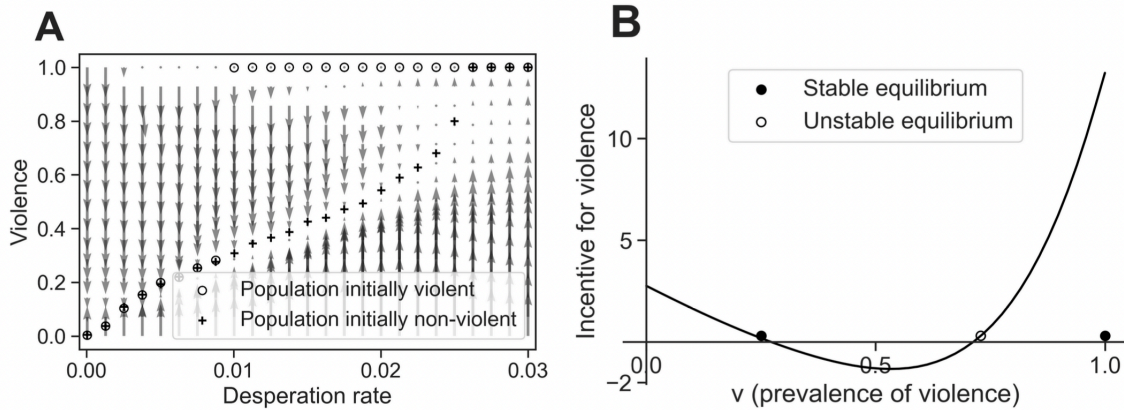


Figure 2.4: (Panel A) Vector field representation of the model results. We obtain the vector field numerically, by initialising the population with a certain proportion of violent agents and a certain desperation rate. We let agents update their strategies and the arrows represent how much the proportion of violent agents has changed. For intermediate desperation rates, the vector field reveals a bistable system, explained by figure 2.4B. (Panel B) Incentives for violence, defined as the mean difference in payoffs between the violent and submissive strategies, obtained with a fixed desperation rate, varying the prevalence of violence. The incentive for violence first decreases as the risk of actually fighting increases, then explodes as the costs of exploitation concentrate on the rare non-violent agents.

2.5 Discussion

Violence rates vary considerably throughout time and space, in association with economic deprivation and inequality, and can have a persistent character. We have shown that a single mechanism can account for these three empirical observations. We therefore offer a candidate explanation, proving that some conditions are sufficient to generate a particular phenomenon. The innovation in this work resides in demonstrating the consequences that flow from making two original assumptions, namely (i) that agents have a desperation threshold and (ii) that violence sends a ‘toughness signal’ that reduces the risk of being exploited. Assumption (i) has been explored in a previous model to explain property crime, but not in relation with violence. Assumption (ii) has been explored earlier in game theoretical models of conflict (Johnstone, 2001; Johnstone & Bshary, 2004; McElreath, 2003), but not yet in relation to deprivation and inequality.

2.5.1 Desperation triggers high violence

The assumption that agents have a desperation threshold has important consequences. First, it triggers risk-proneness below the threshold (de Courson & Nettle, 2021). Intuitively, desperate individuals have ‘little to lose’: if their gamble succeeds, they lift their head above water, and if it fails, it makes little difference. As we assume that stealing entails the highest variance, ‘desperate’ agents are likely to exploit. This result holds as long as individuals possess such a threshold. The ‘threshold’ idea can have several interpretations, starvation being the most obvious, but perhaps not the most relevant for the inhabitants of industrialised countries.

Therefore, a larger proportion of desperate agents – as a consequence of either poverty or inequality (figure 2.2) – leads to more violence (figure 2.1), as non-desperate individuals try to reduce the risk of exploitation. The exploitation that stems from being below a desperation threshold is largely insensitive to the magnitude of punishment (de Courson & Nettle, 2021), in line with empirical evidence (Nagin, 2013). As individuals scramble to get back above the threshold, desperate individuals care more about the maximum payoff than the expected payoff of their strategy. In our model, this means that desperate agents continue to steal even when facing a high risk of violence. As a consequence, violence acts as a ‘deflector’ rather than a deterrent: it will not prevent exploitation, but might make the offender shift to a different, non-violent target.

This ‘deflector’ property of violence fundamentally influences the results of the model. In conventional rational choice models (Becker, 1968), violence deters stealing by increasing its potential cost. In such models, violence acts as a ‘thermostat’, dampening variation in rates of property crime – a sort of homeostasis. This prediction is at odds with empirical evidence showing massive variation in rates of violence across space and time (Glaeser et al., 1995). Instead our model predicts, due to the desperation threshold, that high rates of property crime can persist despite high rates of violence, which is more consistent with the empirical record (Quick et al., 2018).

Despite the common intuition that ‘violence begets violence’, standard game theory models actually predict the reverse. For instance, the Hawk-Dove model (Maynard Smith & Parker, 1976) finds violence to be a negatively frequency-dependent strategy: as violence becomes more common, a violent individual is more likely to meet another violent individual and to get into a costly fight. Put differently, every increase in violence diminishes its appeal, which stifles its spread. For that reason, the Hawk-Dove model predicts that the population only reaches a pure equilibrium where all individuals adopt a violent strategy when the cost of losing a fight is smaller than the cost of the resource at play (Maynard Smith & Parker, 1976). Again, this suggests that violence should display little variation between communities: the costs and benefits might vary, but the negative frequency dependence should homogenise violence rates. This mechanism – violence becoming more costly as it spreads – also operates in our model. However, the assumption that violence confers a ‘toughness reputation’ counteracts this dynamic. As being violent deflects property crime on non-violent agents, the spread of violence also makes non-violence more costly. The frequency dependence reverses for high enough levels of violence and turns into a positive feedback, whereby violence actually begets violence. Which one of these dominates depends heavily on the prevalence of exploitation, which in turn depends on the proportion of desperate individuals. Thus, our model recovers the potential for rates of violence to vary sharply in a way that depends notably on socioeconomic deprivation.

In some regions of the parameter space, high- and low-violence configurations can simultaneously be stable equilibria in our model. Depending on the starting point, the population can end up in either of the two equilibria: if violence is low, then it stays low due to the negative frequency dependence; if violence is high, it stays high due to the positive feedback. If violence is rare, it is not worth sending a violent signal as the

risk of victimisation is diluted in the population, whereas if the vast majority of individuals are violent, a rare non-violent agent will bear the brunt of victimisation and suffer untenable costs. To understand this result more intuitively, one can think of bike locks, which play a protective role analogous to violence in our model. When leaving your bike among dozens of unlocked bikes, it may not be necessary to lock it, as the risk of it being stolen is diluted among all the bikes. However, if all the bikes around are locked, a stealer passing by would likely steal the only unlocked bike. Similarly, the risk of being exploited might be low enough in the low-violence equilibrium for the cost of violence to be too high to incur. In the high-violence equilibrium, however, any submissive deviant will inevitably concentrate the risks of exploitation.

This situation is analogous to coordination games, where positive frequency dependence generates multiple equilibria – for instance left- and right-hand driving. But whereas collective wellbeing is roughly equal whether cars drive left- or right-wing, in our case, settling in a high- or low-violence equilibrium is very consequential. An earlier evolutionary game theory model of toughness signals also reports the possible coexistence of two equilibria with very different levels of aggression (Johnstone & Bshary, 2004). The authors conclude that a population is unlikely to persist in the high-violence equilibrium, as mean fitness is lower than in the low-violence one. Our model does not include an equilibrium selection process. But in the human case and for the relatively short timescales we are interested in, extending only to a few generations, it seems plausible that a community could be trapped in such a detrimental equilibrium, and that such an equilibrium can therefore be empirically relevant.

This bistability generates a hysteresis effect: violence rates do not only depend on the current economic conditions, but also on their history. Concretely, a neighbourhood can be more violent than an equally rich neighbourhood because it was poorer and more violent in the past. Thus, violence can persist despite some economic improvement. This result aligns with the ‘enduring neighbourhood effect’ (Sampson, 2012) of violence, and thus offers an alternative explanation for it – not necessarily incompatible with the prevailing approach, the social efficacy theory (Sampson, 1997).

2.5.2 Relative poverty or absolute poverty?

In our model, poverty and inequality only increase violence through the proportion of desperate agents. Thus, the effect is driven only by absolute poverty, and not by relative poverty. Concretely, our model predicts that making the richest individuals richer with no impact on the poorest does not increase violence. This might seem at odds with several empirical findings (Blau & Blau, 1982; Daly, 2017; Kelly, 2000; M. Wilson & Daly, 1997) and theories (Daly & Wilson, 2001; Krupp & Cook, 2018; Pickett & Wilkinson, 2015; R. Wilkinson, 2004). For instance, comparing the neighbourhoods of Chicago, Daly and Wilson report that when controlling for economic inequality, median household income did not predict homicide rate in Chicago neighbourhoods (M. Wilson & Daly, 1997). However, our model’s prediction is a consequence of our assumption that the desperation threshold is exogenously fixed across communities. Instead, the threshold itself may increase with affluence and with inequality. For example, the existence of inequality has been found experimentally to increase individuals’ “perceived needs” (Payne et al., 2017). In our model, if the threshold was assumed to be proportional to the median resource level (as the poverty line is defined in the European Union), then any change in poverty through the parameter would affect equally both agents’ level of resources and the threshold, without consequence. In this setting, the level of violence would only depend on inequality (σ). One could also imagine that an exploited agent loses a proportion of his wealth instead of a fixed amount of resources. In this setting, the presence of very well-off individuals could act as an incentive for exploitation,

which could in turn create a stronger need for protection among these individuals.

2.5.3 Future directions

Our model analyses violence as a binary decision, where choosing ‘violence’ essentially means being ready to fight when exploited, which yields a ‘toughness reputation’, known by all of the other agents. This assumption is restrictive in two respects that we could explore in the future. First, we could allow agents to fine-tune their level of violence. This would require specifying how the fitness costs of violence relate to the level of violence and how exploiters choose their targets in this new context. It would be interesting to test, for instance, whether the hysteresis effect holds in this case, and to observe if the insensitivity to deterrence induced by the desperation threshold triggers arms race dynamics through an ‘inflation of toughness’.

Second, our assumptions implicitly entail that an agent simply needs to be ready to fight so that every other agent treats him as tough, even if she has never fought. In our simulations, the majority of violent agents reap the benefits of violence without actually paying the costs, as fights rarely occur. We chose here to directly specify how the signal changes the receiver’s behaviour, in order to focus on the population dynamics. An alternative assumption would be to only attribute a toughness reputation to an agent after she actually fought. This would make the model more complex, but would enable it to study another question: under what conditions should an agent be violent without material necessity, for pure reputational reasons? This could shed light on Anderson’s observation that, in the deprived communities he studies, “there are always people around looking for a fight”, a situation he attributes to “campaigning for respect” (Anderson, 2000).

Finally, our model also generates novel testable predictions and can guide empirical data collection. For instance, we predict individuals to be violent when they live in a deprived neighbourhood, whether they are themselves materially deprived or not (figure 2.1). This suggests that an individual’s attitude toward violence would be better predicted by the economic situation of their neighbourhood, rather than by their own personal socio-economic status, whereas the reverse would be true for property crime. It is reminiscent of Anderson’s observation that in deprived neighbourhoods, relatively better-off “decent families” reluctantly teach the “code of the streets” to their children, as they consider a toughness reputation as necessary to navigate this social environment (Anderson, 2000). Cohen and Nisbett also found that geographical origin (North or South) determined whether individuals endorsed “culture of honor” values, regardless of their current circumstances (D. Cohen & Nisbett, 2016).

2.5.4 Implications for policy and intervention

We discuss implications of our model for public policy and intervention. First, our model predicts that helping the poorest individuals to get back above the threshold of economic deprivation can have a ripple effect. Not only would such help reduce the probability that desperate individuals choose to steal, violence would also diminish among the rest of the population, for whom there has been no change in resource levels. In our model, the dire poverty of some has an emergent effect on the population: in a world where some individuals are ready to do anything to get their head above water, everyone must take protective steps like violence. By eliminating desperation, one might then improve the welfare of all members of the group, even those who do not benefit from the policy.

The second implication is related to the hysteresis effect our model produces. For a community trapped in a high violence equilibrium (figure 2.3, upper branch), the effect of an economic intervention would be a step function (figure 2.3, top arrows). Minor interventions would have no effect, then at some point, a sufficiently

large intervention would cause a phase transition to a low violence equilibrium, and have a massive impact (figure 2.3, vertical downward pointing arrow). Similarly, interventions on violence (i.e. lowering violence without changing the desperation rate) would need to be sufficiently large to have an effect. In figure 2.4B, it would be necessary to pass the unstable equilibrium to reach the basin of attraction of the low-violence equilibrium, otherwise violence goes back to maximum when the intervention ends. Therefore, our model suggests a non-linear relationship between intervention magnitude and resulting change. It also calls into question the use of linear models to test the efficiency of these interventions. An empirical researcher might for instance find a null effect of economic support on violence if she only looks at a limited intervention, whereas more of this same intervention could have a tremendous positive impact.

To conclude, we presented a simple model combining individual optimal decisions and population simulations. We make two original assumptions: agents have a desperation threshold, and violence serves as a ‘toughness signal’. We show that their combination is able to explain three cornerstone empirical findings: the large variation of violence rates between neighbourhoods or communities; the effect of poverty and inequality on violence; and the persistence of violence across time.

Funding and acknowledgements

The authors report no conflict of interest. WEF’s contributions have been supported by the Dutch Research Council (V1.Vidi.195.130) and the James S. McDonnell Foundation (<https://www.jsmf.org/grants/20170007/>). DN’s contribution was supported by ANR grant ANR-21-CE28-0009. JLVG’s contribution was supported by a Consolidator Grant from the European Research Council [Grant Number 772911– CRIME-TIME]. The authors thank Julien Lie, Karthik Panchanathan, Daniel Nagin, Ulysse Klatzmann, Sebastian Kübel, Laura Pick, Daniel B. Krupp and two anonymous reviewers for their precious help.

Data transparency and reproducibility

The Python code of the model and of the figures can be found in this repository: https://github.com/regicid/model_deprivation_violence. An online notebook to manipulate the model online without installing Python is available here: https://colab.research.google.com/drive/1wf3KBd95YO_WTluGztaR-8l-1zOD-eoo?usp=sharing.

2.6 Appendix

2.6.1 The model

Notation used in the model

Symbol	Meaning	Typical values
s_t	State (i.e. resource level) at time t	

Symbol	Meaning	Typical values
r	Correlation between s_t and $s_t + 1$	0.99
σ	Inequality (standard deviation of states)	
μ	Affluence (average state)	
p	Prevalence of exploitation	
v	Prevalence of violence	
n	Number of possible targets	10
N	Population size	
ω	Fitness reduction imposed below the desperation threshold	.1
λ	Fitness reduction imposed when losing a fight	.1
β	Stake of exploitation	5
γ	Probability of being caught	1/3
π	Cost of being caught	10
m	Probability of violent misfire	.01
T	Number of SDP iterations	100

2.6.1.1 States

An agent is defined by a dynamic state variable $s_t \in [-50, 50]$, which represents his level of resources at time t . It fluctuates through time as the result of actions taken and random perturbations. As adding these random perturbations to resources levels increases inequality of state in the population by ‘diffusion’, we add a reversion to the mean set to precisely offset this effect. This gives a stationary AR(1) process, which maintains on average the mean state $\mu = \sum_i \frac{s^i}{N}$ and the variance $\sigma^2 = \sum_i \frac{(s^i - \mu)^2}{N}$ of the population (for a formal proof, see (Bateson & Nettle, 2017)):

$$s_{t+1} = (1 - r) \cdot s_t + r \cdot \epsilon$$

$$\epsilon \sim \mathcal{N}\left(\mu, \frac{1 - r^2}{(1 - r)^2} \cdot \sigma^2\right)$$

The term $r \in [0, 1]$ is the correlation between an agent’s current and subsequent level of resources. This process is a shuffling of agents’ state, but is not ‘redistributive’: population inequality is left unchanged. The state space $([-50, 50])$ is discretized with 1001 steps of size 0.1.

2.6.1.2 Actions description

Agents choose among three strategies:

- ‘Exploitation’, i.e. trying to take β resources from another agent. An exploiter has an exogenous probability γ of being caught by the police, in which case he loses π resources. The exploiter chooses whom to exploit in a sample of n individuals, drawn randomly from the population. Among them, he picks uniformly among the least violent (see below).
- ‘Violence’, i.e. fighting when another agent tries to exploit you. The winner of the fight is chosen by a coin toss and gets or keeps the disputed β resources. Choosing this strategy also send a toughness signal to exploiters: when choosing among the n possible targets (see above), exploiters know which agents are violent and pick, if there are any, among submissive agents (defined below).

Violent agents also have a small probability m of attacking when not exploited (a ‘misfire’), which triggers a fight if the attacked agent is also violent.

- ‘Submission’: never fight.

2.6.1.3 Consequences of actions

Formally, decisions are defined by their probabilistic consequences on state. The expected consequences depend on the proportions of exploiters and of violent agents in the population. We denote these respectively p and v .

- An exploiter has a probability γ of getting caught, which costs π units of resources. If not caught (probability: $1 - \gamma$), he gets β resources if there is a submissive individual in the sample, or if he wins the fight against the (necessarily) violent target. The probability is thus: $(1 - \gamma)(1 - \frac{v^n}{2})$. He loses a fight with probability $(1 - \gamma)\frac{v^n}{2}$, which slashes the fitness by a proportion λ .
- Every exploiter draws a sample of n individuals, so a non-exploiter has a n/N chance of being in such a group. In a round (defined by a process in which every exploiter makes a move, successful or not), the number of such interactions is distributed as a binomial distribution $\mathcal{B}(Np, \frac{n}{N})$. p being typically very small, we neglect for simplicity the unlikely possibility ($\sim np^2$) of multiple encounters and model it as a binary event with probability pn , the expected number of encounter with an exploiter.
- If possible, the exploiter targets one of the submissive individuals. If there is none, then he targets a violent individual. So:

- (i) A violent individual will only be picked if every other agent in the sample is also violent (probability: v^{n-1}), and if he is picked among them (probability: $\frac{1}{n}$). So, for a violent individual,

$$P(\text{exploitation}|\text{encounter}) = \frac{v^{n-1}}{n}$$

- (ii) When there are submissive agents among the n available targets, the exploiter picks uniformly among them. For a submissive agent, the probability of being exploited is thus $\frac{1}{K+1}$, K being the number of others submissive individuals in the sample, distributed as a binomial with v and $n - 1$ as parameters. The law of total probability gives:

$$P(\text{exploitation}|\text{encounter}) = \sum_{i=0}^{n-1} \frac{P(K = i)}{i + 1}$$

$$\begin{aligned}
&= \sum_{i=0}^{n-1} \binom{n-1}{i} (1-v)^i v^{n-1-i} = \sum_{i=0}^{n-1} \frac{(n-1)!(1-v)^i v^{n-1-i}}{(i+1)! (n-1-i)!} \\
&= \sum_{k=1}^n \frac{(n-1)!}{k!(n-k)!} (1-v)^{v-1} v^{n-k} \\
&= \frac{1}{n(1-v)} \left[\sum_{k=0}^n \binom{k}{n} v^k (1-v)^{n-k} - v^n \right] \\
&= \frac{[(1+(1-v))^n - v^n]}{n(1-v)} = \frac{1-v^n}{n(1-v)}
\end{aligned}$$

- When not exploited (probability: $(1 - pv^{n-1})$), a violent individual attacks anyway with probability m . If the attacked agent is also violent (which happens with probability v), this escalates into a fight, which he loses with probability $\frac{1}{2}$. If the attacked agent is submissive, nothing happens. All in all, a violent individual has a probability $\frac{mv}{2}(1 - pv^{n-1})$ of losing a fight, which slashes his fitness by a proportion λ .

2.6.1.4 Fitness function and optimisation

Agents have an horizon of T time steps. Their fitness is defined by their terminal level of resources (s_T), reduced by penalties in two cases:

- Each round t where $s_t < 0$ (i.e. below the ‘desperation threshold’) reduces fitness by a proportion ω
- Each lost fight reduces fitness by a proportion λ , to represent possible health sequelae.

The fitness is thus:

$$s_T \cdot (1 - \omega)^{\#(i / s_i < 0)} \cdot (1 - \lambda)^{\#(\text{lost fights})}$$

Agents compute the expected payoff of the three strategies by stochastic dynamic programming (SDP) ((Houston et al., 1999; Mangel & Clark, 1988)). We use dynamic rather than static optimisation so that agents take into account not only the probability of falling below the desperation threshold in the next time step, but also at any point in the future. As we increase T , the algorithm converges to an optimal long-term policy. We choose a value of T large enough (100) to obtain that convergence.

The model is not a full game theory model. There are strategic interactions (agents’ payoff depend on others’ strategies). Agents take into account the distribution of strategies in the population but not the consequences of their decision on others. As we simulate large and well-mixed populations, we can safely consider this effect as negligible. We also assume, for tractability, that agents optimise as if p and v stayed constant in the future - which is true at equilibrium.

2.6.1.5 Population simulations

We simulate populations of N agents. Initial states are drawn from a Gaussian distribution with mean μ and variance σ^2 . To decide their strategies, agents need to measure the frequency of violence and exploitation in the population (p and v), which raises a chicken-and-egg problem. To circumvent it, we first initialise strategies by letting agents choose the optimal strategy for $p = v = 0$. We then exogenously fix v (the proportion of violent agents) in the population either to 0 or to 1 by assigning either the submissive or the violent strategy to non-exploiting agents. This allows us to test whether the end results change depending if we start from an initially violent or non-violent population (figure 2.3 & 2.4 in the paper).

We then let, at each time step, 10% of the population update their strategies. This gradual updating allows the equilibrium to be smoothly reached, and avoids large oscillations. When updating their strategies, agents measure p and v in the population and choose the optimal strategy for their state by the aforementioned SDP algorithm. We run the simulations for 100 time steps, so that the population can reach an equilibrium in the distribution of strategies. During the simulations, we let the resource levels of the agents fluctuate following the stochastic AR(1) process described in Section 2.6.1.1, but we neglect the consequences of agents' actions on the overall population affluence and inequality. In other words, we assume that individual behaviours have negligible consequences on the economic context. We make this choice to preserve the Gaussian distribution we started from. Including the effect of agents' actions on their state often prevents reaching an equilibrium. In particular, agents below the desperation threshold choose exploitation and have thus a negative expected payoff, which leads them to 'drown'. Hence an ever-increasing value of p , as more agents fall below the desperation threshold by stochasticity and are then stuck below. The choice to neglect actions' economic consequences in population simulations reduces their realism but not their interest: the model aims to explore how the economic context (the distribution of resource levels), treated exogenously, determines behaviour. Exploring potential feedback between behaviours and economic context, or tracking individual-level life course trajectories, is beyond our present scope.

2.6.2 Supplementary results

2.6.2.1 Fitness function

The SDP algorithm computes implicitly a "value function" comparing the expected fitness of all actions. As the backward induction advances, its shape tends to stabilise in what we can interpret as a long-term fitness function. The typical shape is shown in Fig. 2.5 and allows to visualise the risk-preferences of the agents. The desperation threshold produces a S-shape, with risk-proneness below and risk-aversion above. This explains why agents tend to exploit below the threshold as it is the strategy associated with the highest variance. It also explains why, in Fig. 2.1 of the main manuscript, agents are especially willing to fight when close to the threshold, where the curve is particularly concave, so where the agents are particularly risk averse. Counter-intuitively, our assumptions make the violent strategy the most risk-averse one: it reduces the risk of exploitation, and the risk of fighting has direct consequences on fitness (a measure where agents are by definition risk-neutral, as it is the maximand). This result might alter or reverse if we instead assumed that losing a fight had consequences on resource levels.




Figure 2.5: Long-term expected fitness depending on the level of resources

2.6.2.2 Sensitivity analysis

In the paper, we use fixed values for most of the parameters defining the rules of the game. In this section, we explore how the model reacts when we vary these parameters and we explore to what extent the pattern we report in the paper are robust to such changes.

2.6.2.3 Varying the costs of exploitation

When exploiting, an agent faces an exogenous risk of punishment, costing π units of resources with probability γ . In a previous model where the same hypothesis was used (de Courson & Nettle, 2021), we found that varying the severity of punishment π was not deterring desperate agents from exploiting, as the fitness function (similar to Fig. 2.5) is very concave below the threshold and quite flat in the lower values). As violence is, in the current model, a response to the risk of exploitation, we would expect π to have no impact on violence. We indeed find no clear effect (Fig. 2.6).




Figure 2.6: Effect of π on the mapping between desperation rate and the proportion of violent agents at equilibrium, starting from a non-violent population.

We can also vary the probability of punishment γ (Fig. 2.7). In de Courson & Nettle (2021), we found a small deterrent effect for this same parameter. Here, we reproduce this finding: the higher γ , the less violence, as some desperate agents (the closest to the threshold) choose not to exploit.




Figure 2.7: Effect of γ on the mapping between desperation rate and the proportion of violent agents at equilibrium, starting from a non-violent population.

2.6.2.4 Varying the costs of violence

When adopting the violent strategy, an agent fights when he is exploited, and also by “mistake” with probability m . When he loses a fight, his fitness is reduced by a proportion λ ¹. The parameters λ and m thus control the costs of violence. Increasing these two parameters strongly reduces violence (Fig. 2.8 & 2.9). Contrary to exploitation (and predictably, as it is not a consequence of desperation), the violent strategy is thus sensitive to its costs.

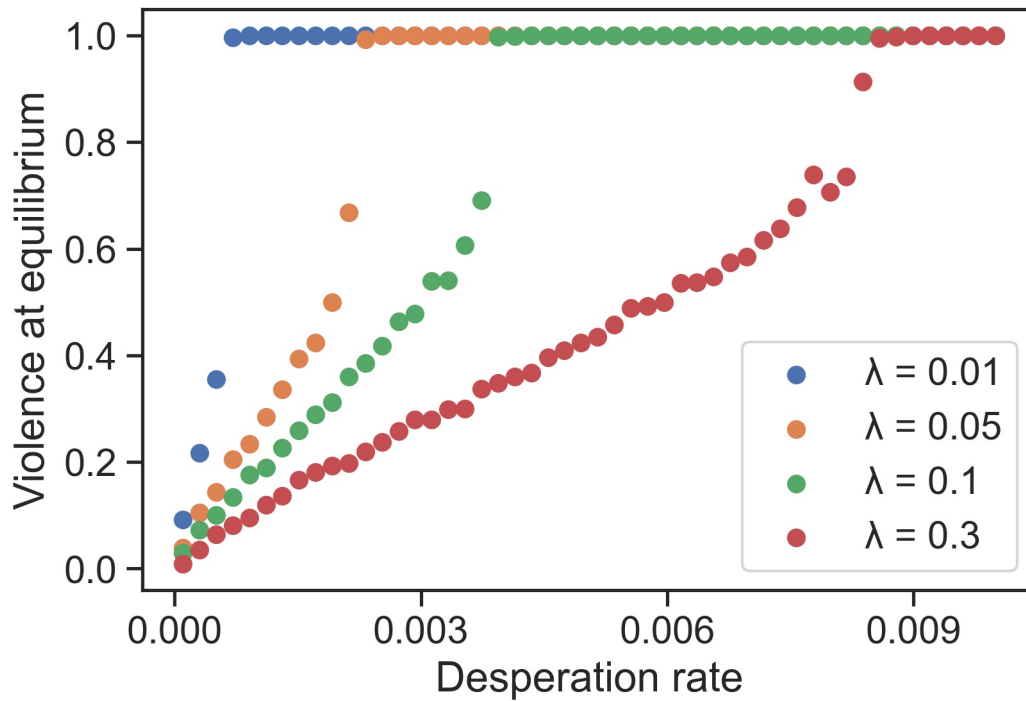


Figure 2.8: Effect of λ on the mapping between desperation rate and the proportion of violent agents at equilibrium, starting from a non-violent population.

2.6.2.5 Hysteresis effect depends on n and r

In the model, exploiters choose their target among n randomly drawn individuals, and choose if possible a submissive one. The parameter n thus controls the “toughness reputation” efficiency: if $n = 1$, a violent agent is as likely to be exploited as a submissive one, whereas if $n \rightarrow \infty$, an agent will never be exploited as long as there are submissive individuals in the population.

¹Increasing ω would have the same effect as decreasing λ . The model’s outcome depends on the relative importance of these two parameters, which determines whether it is more important for an agent to be below the threshold or to avoid losing fights. For that reason, we explore the consequences of varying λ but not ω .




Figure 2.9: Effect of m on the mapping between desperation rate and the proportion of violent agents at equilibrium, starting from a non-violent population.

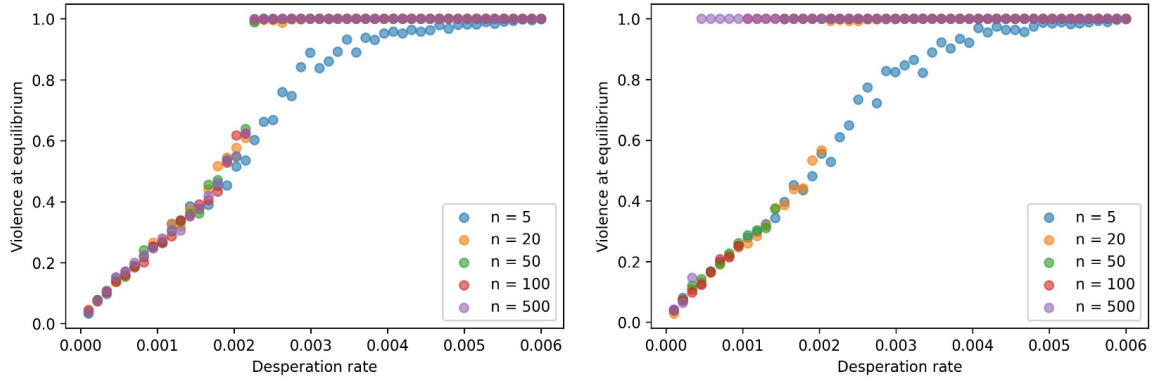


Figure 2.10: Effect of n on the mapping between desperation rate and the proportion of violent agents at equilibrium, starting from a non-violent (left) or violent (right) population.

Modifying the value of n has two effects on the model's dynamics.

First, the model typically produce a S-curve between desperation rate and the proportion of violent individuals, with a concave then a convex phase (Fig. 2.3 in the main manuscript). The value of n does not affect the concave phase, but makes the convex one steeper (Fig. 2.9). The S-shape can be interpreted in terms of frequency dependence of the violent strategy: violence is negatively frequency-dependent at low levels, as in the Hawk-Dove model, and positively frequency-dependent at high levels, because violence deflects exploitation on the rare submissive individuals (see Fig. 2.4b in the main manuscript). n controls this deflection effect: if $v \rightarrow 1$, then the risk of being exploited converges to p for a violent agent and to pn for a submissive one (see Sec. 2.6.1.3). It is thus logical to observe that n affects the convex phase. For low v (the concave phase) violent agents are almost sure not to be exploited, so the value of n makes little difference. This phase actually depends on the parameters m and λ , which control how frequently rare violent individuals will fight and how costly it is (see Fig. 2.8 and 2.9). Second, and also because it controls the deflection effect, n determines if we obtain one or two equilibrium, i.e. the hysteresis effect. In Fig. 2.11, we see that the initial level of violence has no impact on the end state for $n = 5$, and that the higher n , the lower the desperation rate needed to obtain two equilibria.

Finally, the results depend on the parameter r , the auto-correlation between an agent's present and future states defined in Sec. 2.6.1.1. The lower r , the more gently violence increases with the desperation rate. This effect comes from the fact that a lower r makes the decision to exploit sensitive to the level of violence: when $v \rightarrow 1$, desperate agents often choose not to exploit, and violence plays thus a regulating role. To understand this result, we can interpret r as a measure of how serious "desperation" is. A low r means that regression to the mean plays a strong role, so that being below the threshold is likely to be an episodic moment rather than a persistent situation. The correlation between s_t and s_{t+x} is r^x , which converges exponentially quickly to zero. In other words, a low value of r implies that a desperate agent can expect to be back above the threshold in the near future without taking any risk.

In this case, an agent may have a higher fitness by waiting for the tide than by taking a gamble which has more chances to fail than to succeed. Interestingly, this effect is offset by increasing the stakes of exploitation β and π (not shown). Desperate agents are still very risk prone, but the risks they are offered are too small to offer a better option as one needs several successful gambles in a row to succeed, which is very unlikely.




Figure 2.II: Effect of r on the mapping between desperation rate and the proportion of violent agents at equilibrium, starting from a non-violent population.

For that reason, a low value of r can suppress the hysteresis effect: in Fig. 2.10, $n = 100$ and $n = 500$ produce no hysteresis contrarily to Fig. 2.10. High-violence reduces exploitation, which destabilises the high-violence equilibrium. The hysteresis effect still occurs with a higher value of n (Fig. 2.12). Thus, the hysteresis result holds for a high enough value of r and n . In other words, it holds if toughness signals are efficient enough and desperation persistent enough.




Figure 2.12: Model results with a lower value of r (.9), starting from a violent population.