

The Onset, Spread, and Prevention of Mass Atrocities: Perspectives from Network Models

Charles H. Anderton* and Jurgen Brauer†

HiCN Working Paper 284

November 2018

JEL Codes: B55, D74, D85, D91

Acknowledgements: A version of this paper was first read at an invited lecture at Washington & Lee University, Lexington, VA, USA, on April 6, 2018. Another version was read as an invited lecture at Chulalongkorn University, Bangkok, Thailand, on June 29, 2018. We thank attendants for comments.

* Professor of Economics, Department of Economics and Accounting, College of the Holy Cross, Worcester, MA, USA

† Emeritus Professor of Economics, Hull College of Business, Augusta University, Augusta, GA, USA, and Visiting Professor of Economics, EBA Program, Faculty of Economics, Chulalongkorn University, Bangkok, Thailand

1. Introduction

History is replete with cases of intentional, systematic, and substantial destruction of noncombant civilians by powerful elites including dictatorial and nondictatorial governments and nonstate groups (e.g., rebel, militia, and terrorist organizations). In just the last 120 years, some 200 mass atrocities have occurred in which at least 100 million civilians have been intentionally killed (and likely at least half of a billion people killed, injured, or traumatized) by political leaders and their fellow perpetrators (Anderton and Brauer, 2018). A substantial literature is now emerging on economic aspects of mass atrocities to complement the already well-developed literatures offered in other social sciences and in the humanities (Anderton and Brauer, 2016, 2018). Furthermore, theories and models of social and economic networks in particular have burst on the scene in the last quarter century and they are updating economists' understanding of a wide range of economic activities including peer effects in education, hub and spoke transportation systems, job search and employment outcomes, the diffusion of information, innovations, and ideas, the roles of interpersonal relationships in market outcomes, and the behaviors of criminal groups including illicit drug trading and human trafficking organizations (Jackson, 2008; Easley and Kleinberg, 2010; Bramoullé, Galeotti, and Rogers, 2016). Moreover, formal networking models are now being developed to analyze conflict on networks including wars and terrorism (Cunningham, 1985; Baccara and Bar-Isaac, 2008; Franke and Öztürk, 2009; Maoz, 2011; Enders and Sandler, 2012; Goyal and Vigier, 2014; Dziubinski, Goyal, and Vigier, 2016; Acemoglu, Malekian, and Ozdaglar, 2016; Zech and Gabbay, 2016; König, Rohner, Thoenig, and Zilibotti, 2017; Scaife, 2017; Terrorism Network Project, 2018). Surprisingly, however, few formal networking models have been applied to mass atrocities. Such applications are warranted because mass atrocities are orchestrated by architects over networks of

perpetrators, play out in locales based upon networks of relationships among perpetrators and resisters, and involve networks between and within both victims and potential rescuers.

Network theory has shown economists and other social scientists that one cannot simply add up the micro actions of individual agents to surmise a macro outcome. Keynes already taught us this with the Paradox of Thrift: When all agents in a macroeconomy attempt to save more, consumption and employment could fall, leading to less savings overall. Neither can one necessarily surmise the underlying micro preferences of individuals from any observed macro, system-wide outcome (Schelling, 1971, 1978; Granovetter, 1978). Much depends on what the network literature calls the aggregation function.

Most importantly for our purposes, network theory can explain certain puzzles associated with mass atrocities. First, while social psychologists have long known that it can be frightfully easy to recruit relatively small numbers of people to commit atrocious acts (e.g., Roth, 2010), how can one explain the more puzzling *mass* participation in *mass* atrocity? Network theory can explain this scaling up. Second, when a mass atrocity such as a genocide breaks out, why can it spread like wildfire in location A yet stay contained in location B even though the individuals in the two locations are essentially the same (even identical) in their initial perspectives on what should happen to the outgroup? Network theory can provide insights on the very uneven spread of mass atrocities at the grassroots level. A third puzzle is this: Why do some individuals in mass atrocities “flip”, i.e., change from perpetrator to rescuer (as did Oskar Schindler) or from rescuer to perpetrator? Fourth, seemingly even more puzzling is the observation in mass atrocity contexts that some individuals rescue and perpetrate *at the same time* (Campbell, 2010; Donà, 2018). Fifth, why do some people who are strongly opposed to acts of mass atrocity nonetheless help commit them (Browning, 2004) and thereby help perpetuate an atrocity norm (Michaeli and

Spiro, 2017)? Sixth, why do some perpetrators, after the fact, display what is almost certainly genuine remorse and, sometimes, sincere candidness that they are shocked at what they did (Gobodo-Madikizela, 2002; Staub, 2005; Wilson, 2010)? Seventh, what explains the “silence of the majority” in many mass atrocity cases? Social network theory provides compelling insights into all these aspects, perspectives, and behaviors without appealing to wildly changing or bizarre preferences or circumstances (e.g., Schelling, 1971; Granovetter, 1978; Gintis, 2009).

At the core of many mass atrocities lies a struggle that plays out *within the in-group* (Kleinschmidt, 1972; Straus, 2006; Adalian, 2009; Lemarchand, 2009; Ray, 2017) over whether the in-group will seek to eliminate some or all of the out-group. If in-group extremists “win” this struggle, mass atrocity commences. This struggle that transpires on the threshold of mass atrocity is a diffusion process. Either atrocity acceptance diffuses to a relatively high level within the in-group beyond which a threshold is crossed or else atrocity acceptance is nullified. To model this preliminary stage of mass atrocity, we begin with the Bass (1969) model of diffusion and then add network aspects of diffusion to this starting point. Assuming that mass atrocity commences, we then offer network modeling of such commencement. Specifically, we offer a stylized linear quadratic model (LQM) of mass atrocity propagation over geographic space along networked “tentacles” of destruction. Among the model’s key results is the dramatic increase in destruction of the out-group that is “afforded” by a “network production function” (which is a comparable concept to a classic production function in standard economic theory). Finally, we add to the network model the contending actions of networked resisters. In so doing, we evaluate certain policies that can degrade the atrocity-perpetrating network, but we also show how a network context can cause certain prevention efforts to backfire and do more harm than good to afflicted civilians.

2. Models of Diffusion among a Population

2.1 The Bass Model of Product Diffusion

2.1.1 Difference Equation and Diffusion Curves

The notion of contagion or diffusion is the spreading of “something” among a population. The “something” might be a disease (e.g., AIDS), an idea (e.g., love your enemy) or opinion (e.g., candidate so-and-so is corrupt), information (e.g., there is a great movie playing at the theater), or, in our context, the “disease” of atrocity acceptance among people within the in-group. The Bass (1969) model is a starting point for analyzing diffusion of products and fashions in the marketing literature. Let $F(t)$ be the fraction of some population that has accepted the product at time t . Either a person has accepted the product (1) or not (0), so the states of the model are binary. Let ρ be the constant rate of *spontaneous adoption* of a product, sometimes referred to in the marketing literature as “early adopters.” In our context, this may refer to *instigators* (or *true believers*), i.e., people who make an adoption decision independent of social, or peer, effects on adoption. Jackson (2008, p. 187) notes that spontaneous adoption can be interpreted as a response to outside stimuli such as advertising. Further, let β be the constant rate of *imitation*. This captures peer effects in adopting the product based on social relations. In peer contexts, adoption of the product can spread by word-of-mouth in personal encounters, over social media, via cell phones, and the like. The basic Bass model then is a purely mechanical, tripartite difference equation that describes the acceptance of a product over time among a population as follows:

$$F(t) = F(t - 1) + \rho[1 - F(t - 1)] + \beta[1 - F(t - 1)]F(t - 1). \quad (1)$$

As mentioned, $F(t)$ on the left-hand side is the fraction of a population that has adopted an item or idea at time t . The first element on the right-hand side of the equals sign is the fraction of

adopters from the last period, $F(t-1)$, i.e., prior adopters, who are assumed to be adopters throughout and thus will continue to be adopters in the current period (t). The next element on the right-hand side of the equals sign captures the rate of spontaneous adoption (ρ) among the fraction of the population that has not yet adopted the product, i.e., $1-F(t-1)$. The third element on the right side of the equals sign captures the peer aspects of adoption (β) from the social interactions between fractions of the population that are nonadopters, $1-F(t-1)$, and adopters, $F(t-1)$, of the product. The parameter β measures the rate of “conversion,” as it were, of nonadopters to adopters. Numerous forms of diffusion behavior are possible in the Bass model depending on the initial condition for F , i.e., $F(0)$, the rate of spontaneous adoption (ρ), and the rate of imitation (β). Figure 1 shows five such diffusion curves. Four of the curves assume $F(0)=0$ (no *true believers* to begin with) while one assumes $F(0)=0.2$ (20% of the population, to begin with).

[Figure 1 about here]

2.1.2 Diffusion of Atrocity Acceptance in the Bass Model

Several intuitive applications of the Bass model to mass atrocities present themselves in Figure 1. First, atrocity architects can manipulate the contagion of atrocity acceptance within the in-group through advertising, i.e., propaganda. Such actions serve to increase the spontaneous rate of adoption (ρ) which pertains to people adopting an idea independent of peer or social influences. (For a review of this aspect see, e.g., Petrova and Yanagizawa-Drott 2016.) For the Rwandan genocide of 1994, Yanagizawa-Drott (2014), for instance, studies the role of hate radio (Radio Télévision Libre des Mille Collines, RTLM, as opposed to Radio Rwanda which did not broadcast inflammatory material) as a coordination device to incite violence. To arrive at statistically credible results, he exploits the quasi-random geographic distribution of hills,

flatlands, and valleys that affects the relative quality of signal reception. Holding other factors constant, communities in reach of a clear line-of-signal indeed showed a statistically significant (and quantitatively large) increase in genocide participation. Moreover, due to spatial spillovers on social interaction networks, so did neighboring communities, suggesting coordination and triggering effects facilitated by hate radio. He estimates that about 10 percent of the participation in the violence is attributable to RTML. This translates to an additional 51,000 perpetrators. (The International Criminal Tribunal for Rwanda later convicted the radio station's founders of instigating genocide.)

Note in Figure 1 the significant increase in the diffusion curve from the one labeled $\rho=0.01, \beta=0$ to the one labeled $\rho=0.04, \beta=0$. At time $t=50$, the former curve has “only” a 39.5 percent adoption rate of the atrocity norm within the in-group, but the latter curve has a rate of 87.0 percent. Clearly, the model suggests that propaganda will be an important tool in the hands of atrocity architects for diffusing the acceptance of atrocity within the in-group.

Second, independent of spontaneous adoption, peer or social effects on adoption of atrocity also greatly matter. Notice the change in adoption curves when moving from $\rho=0.01, \beta=0$ to $\rho=0.01, \beta=0.04$ to $\rho=0.01, \beta=0.1$. At time $t=50$ across these three curves, the rate of adoption rises from 39.5 percent to 69.0 percent to 95.9 percent. Furthermore, the emergence of an S-shape adoption curve *requires* that sufficiently strong peer effects be present. The S-shape of a diffusion curve is important in the marketing literature because it reflects that part of a diffusion process in which acceptance is *increasing at an increasing rate* before it eventually increases at a diminishing rate. Numerous cases studies of mass atrocities reflect a dramatic uptick in mass atrocity acceptance within the in-group generally and in specific locales in which

mass atrocity perpetration takes off suddenly (Browning, 1992, ch. 8; Des Forges, 1999; Fujii, 2009).

Thirdly, the initial fraction of adopters, $F(0)$, determines the height of the launching pad for diffusion. In Figure 1, the difference between the $\rho=0.01$, $\beta=0.1$, $F(0)=0$ and $\rho=0.01$, $\beta=0.1$, $F(0)=0.2$ curves is quite dramatic. At time $t=20$, the former curve's fraction of adoption is 40.5 percent, while the latter's is 75.0 percent. Assuming that $F(0)$ is the fraction of the population predisposed to atrocity acceptance (potential *true believers*), even before there is one, then once atrocity begins they help feed the peer aspects of the Bass model right from the get-go, leading to rapid diffusion in a relatively short time. If time is measured in days, then the top diffusion curve in Figure 1 reaches 50 percent acceptance by day 11. Meanwhile, the lower diffusion curve in the figure reaches 50 percent adoption by day 69. These are dramatically different diffusion scenarios with drastically different third party intervention prospects. If policy efforts are reactive, perhaps even timely by the standards of international policymaking, then the diffusion along the top curve will likely be well on its way before policy intervenes and attempts to make a difference. Such a scenario will tend to be a policy that is too little, too late.

Although we have applied ideas from the Bass model to an overall intra-group struggle over atrocity policy, diffusion processes matter greatly at the grass roots level, too. Hence, we will return to processes of diffusion at local levels once we have developed the model of atrocity tentacles reaching into locales later on in this article.

2.1.3 The Bass Model and Atrocity Prevention

Assume initially that we have nonnegative values for the rate of spontaneous adoption (ρ), the rate of imitation (β), and the initial fraction of adopters ($F(0)$) in the Bass model. Prevention of diffusion of atrocity acceptance in the population requires that the rate of spontaneous adoption

(ρ) and the initial fraction of adopters ($F(0)$) be zero or that the rate of spontaneous adoption (ρ) and the rate of imitation (β) be zero. These are heroic assumptions to expect of a population of relatively large size (such as a village). There will almost always be some people with an exogenous animus toward people from the out-group, implying $\rho > 0$ and $F(0) > 0$. Hence, the model suggests a latent potential for atrocity acceptance if the diffusion process is set into motion by atrocity architects. As Valentino (2004, p. 2) notes, “the minimum level of social support necessary to carry out mass killing has been uncomfortably easy [for leaders] to achieve.” Thus, the model helps explain the first puzzle: how an animosity within a small group toward an out-group can mushroom into mass participation across large segments of the population.

Naturally, if other parties (say third parties) are able to cultivate negative values for ρ and for β (or a sufficiently large negative value on one of the parameters), a process of atrocity rejection can be fostered as shown in Figure 2. This is true, as the figure shows, even for high initial fractions of atrocity adopters, $F(0) \gg 0$. Policies that might make ρ and/or β negative include counter-propaganda campaigns to offset and turn around hate radio, threats of litigation against those who adopt atrocity actions against the out-group, insertion of third party peacekeepers to reduce $F(0)$ and turn ρ and/or β toward negative values. It is also important to note how the ultimate outcomes in the Bass model can vary dramatically based upon a seemingly trivial change in a parameter. For example, the $\rho=0.01$, $\beta=-0.03$ diffusion curve in Figure 2 will eventually reach an outcome in which the whole population is atrocity acceptant ($F=1$) (not shown) as the negative peer effect parameter is not strong enough to overcome the initial true believer parameter. Yet, if a slightly more negative imitation parameter of $\beta=-0.04$ was in place (all else equal), then the $\rho=0.01$, $\beta=-0.04$ diffusion curve in Figure 2 would reach an outcome in

which the whole population rejected atrocity ($F=0$) (not shown). Hence, the ultimate fate of an intra-group struggle in a locale, whether it tips to atrocity acceptance or rejection, can be extremely sensitive to initial conditions and small parameter changes. This helps explain the second puzzle, why mass participation in an atrocity can spread rapidly in one location yet fail to diffuse to or in another.

[Figure 2 about here]

2.1.4 Epidemiological Models of the Spread of Disease and the Bass Model

The literature on the spread of diseases over networks of people, animals, and plants is vast (Jackson, 2008, 185–221; Newman, 2010, 627–73; Lamberson, 2016). Models of epidemics often involve one of four epidemiological models known by their acronyms—*SI*, *SIR*, *SIS*, and *SIRS*—where “*S*” stands for “susceptible,” “*I*” for “infected,” and “*R*” for “recovered.” We leave applications of such models to mass atrocity to future research, but we do so with three caveats. First, the *SI* model from the epidemiology literature is a special case of the Bass model in which the spontaneous adoption parameter (ρ) is set to zero. Second, while there are important examples in mass atrocity contexts of individuals who were susceptible, became infected with mass atrocity perpetration, and then recovered (i.e., the *SIR* model, e.g., Oskar Schindler, and thus addressing the third puzzle), far more common are cases of people becoming habituated and “locked in” (unrecoverable) to atrocity through various social psychological processes such as the “foot in the door phenomenon,” “motivated rationalizations of atrocity actions,” and peer group effects (Waller, 2007). If we assume, then, that most or all of those who become infected stay infected (which puts us in the Bass world), then the fraction of those infected (assuming, for instance, non-negative ρ and β in the Bass model in which at least one of the two parameters is positive) *always goes to one*. This stark outcome is not a realistic description of mass atrocity

contexts in which there are always some who remain immune (i.e., resisters) regardless of how many others in the population become infected (i.e., perpetrators). In short, the third caveat is that the diffusion process in the Bass model needs to be generalized to account for polymorphic outcomes. To do so, we turn to analysis of a generalized S-shaped model of diffusion.

2.2 Generalized S-Shaped Model of Atrocity Diffusion and Prevention

In this subsection, we adopt the Easley and Kleinberg (2010, ch. 17) model of diffusion of a new product in the marketplace and adapt it to the context of diffusion of atrocity acceptance among a population encompassing an in-group. In so doing, we develop a general S-shaped model of diffusion in which the diffusion curves generated by the Bass model are special cases. The S-shaped model also generates insights into atrocity rejection among the in-group population.

2.2.1 Intrinsic Valuation of Atrocity Perpetration

Following Easley and Kleinberg (2010, p. 450), assume each individual from the in-group population has a “name” given as a real number between 0 and 1, that the number of such individuals is finite, and that the total mass of such individuals is 1. The set of individuals with names between 0 and x (where $x \leq 1$) represents the fraction x of the in-group population. Assume each individual x has a spontaneous or intrinsic interest in participating in acts of harm against the out-group as represented by a *reservation value*, $r(x)$, for participation. The higher the reservation value for any given cost of participation, the more likely the individual would prefer to be an atrocity perpetrator rather than an abstainer. Overriding cost considerations, some individuals may have very high reservation values for participation owing to latent hatred of the out-group. Others might have zero reservation values and, all else equal, would not participate no matter how low the cost. Still others might find some intrinsic value in participating in harm against the out-group owing to looting opportunities. Since reservation values can vary across

individuals in the population, heterogeneous preferences are implied. Each individual x 's intrinsic reservation valuation, $r(x)$, is independent of any peer or network effects that might cause the individual to value participation differently (this network element will be introduced shortly).

Assume that an atrocity architect would like to recruit individuals into atrocity perpetration. The potential availability of personnel would be based on an aggregate reservation value function such as shown in Figure 3. This curve (which, for convenience, is drawn as a straight line) can be thought of as the “market demand” curve for participating in atrocity by potential recruits. The demand for atrocity perpetration is given by the reservation valuation function, $r(x)$, for the range of individuals between 0 and 1. The cost of participating in atrocity, c , involves the cost of effort and, for at least some actors, unpleasant side effects from harming people from the out-group. In most (civil) societies, atrocity actions are crimes subject to the costs of prosecution and the possibility of incarceration. In such societies, c would be high and few individuals would participate in hate crimes. In societies in which political leaders seek to foster atrocity, however, the cost of participation might be lowered owing to decriminalization of hate crimes against out-groups. In addition, propaganda might be used to increase the reservation values (demand) associated with atrocity participation. In Figure 3, the combination of the aggregate demand, or society-wide reservation valuation, curve and the cost of atrocity perpetration gives rise to the fraction x of the in-group participating and the fraction $(1-x)$ not participating. If the cost of atrocity perpetration is $c_I=0.25$ as shown in the figure, individuals with names between 0 and $x_I=0.75$ will become perpetrators and individuals with names greater than $x_I=0.75$ will not. (If we draw the reservation valuation curve in the figure such that it turns down and intersects the horizontal axis before reaching individual 1, it would imply that some

individuals would not participate in atrocity even if the cost of doing so is $c=0$.) Note that the implied reservation value function in Figure 3 is $r(x)=1-x$. This, of course, is a special case of a more general down-sloping reservation value function.

[Figure 3 about here]

2.2.2 Network Effects in the Model

The analysis in Figure 3 ignores network effects in which the atrocity participation of one or more members of the population can alter the valuation of atrocity perpetration of others in the population. We now add network benefits from atrocity perpetration to the intrinsic benefits shown in Figure 3. Assume now that individual x 's valuation of atrocity perpetration involves its spontaneous or intrinsic valuation, $r(x)$, and its benefit from having a fraction z of the in-group population on board, with the atrocity architect's aims represented by the value function $f(z)$ in which $f(0)=0$ and $f'(z)>0$. This peer-effect function, f , captures the network benefits available to individuals from in-group participation (for example, benefits from peer support, information flows, ability to rationalize participation when others are participating, and so on). Here z represents individuals' shared expectation of the fraction of the population that will adopt atrocity perpetration. The modified reservation value for adoption of any individual x from the in-group will now be $r(x)f(z)$, i.e., own-benefits times peer-mediated benefits flowing to the individual, where the multiplicative form implies that individuals with high intrinsic value from atrocity perpetration benefit the most from increases in the fraction of the population that participates (Easley and Kleinberg, 2010, p. 455). A *self-fulfilling expectations equilibrium* is one in which the people in the population expect the fraction z to adopt perpetration and the fraction that actually adopts also is z (Easley and Kleinberg, 2010, p. 454). For any given cost of atrocity perpetration, say $c=2$, what value (or values) of z will be an equilibrium in this self-

fulfilling sense? We can immediately establish that $z=0$ will be a self-fulfilling expectations equilibrium. If everybody expects a proportion $z=0$ of the population to perpetrate, then the reservation value for each individual x is $r(x)f(z)=r(x)f(0)=r(x)0=0$. In this case, each individual's valuation for participation is zero and thus less than the cost, $r(x)f(z)=0 < c=2$, so that nobody in the population has an incentive to participate and the shared expectation of $z=0$ indeed is fulfilled (Easley and Kleinberg, 2010, p. 454).

Might there be a positive value (or values) of z that would be a self-fulfilling expectations equilibrium? The answer will depend on the nature of the r and f functions and the cost of atrocity adoption. To answer the question, consider the following special cases for the two functional forms. Suppose as above in Figure 3 that the r function takes the form $r(x)=1-x$. Recall that z is the society-wide expectation of the fraction of the population that will participate. Since everybody (including person z) expects that person with number name z will adopt atrocity, it follows that all individuals with number names $x < z$ will be expected to adopt atrocity. Hence, all individuals from 0 to z will have an intrinsic reservation value of *at least* $r(z)=1-z$. Let the f function of network benefits be $f(z)=12.5z$. Figure 4 plots the aggregate reservation value function $r(z)f(z)=(1-z)(12.5z)$ across various levels of expectations z , assuming that the cost of participation is $c=2$. Given the specific functional form for reservation value, if nobody in the population is expected to participate then, to repeat, there are no network benefits, $f(0)=0$, and thus the reservation value for each person is $r(z)f(z)=r(z)(0)=0$. In this case, each individual's valuation for participation is less than the cost, $r(z)f(z)=0 < c=2$, so nobody in the population has an incentive to participate. Hence, when no one is expected to participate, no one will value participation, and a shared expectation of $z=0$ is a self-fulfilling expectations equilibrium as already established above.

What other self-fulfilling expectations equilibria exist in Figure 4? Suppose the atrocity architects are able to jump over the “hurdle” of $z=0$ expectations of the population and foster an expectation fraction of participation of $z'=0.2$ in Figure 4. The reservation valuation for the person with number name $z'=0.2$ will adopt perpetration because his reservation valuation just meets the cost, $r(z)f(z)=(1-z)(12.5z)=(1-0.2)(12.5)(0.2)=2=c$. Further, all individuals with number names less than 0.2 will also adopt because their reservation valuations will be greater than 2. For example, person 0.1’s reservation value will be $(1-0.1)(12.5)(0.2)=2.25>c=2$. Individuals with number names greater than 0.2 will not adopt perpetration as their reservation valuations will be less than 2. For example, person 0.3’s reservation value will be $(1-0.3)(12.5)(0.2)=1.75<c=2$. Hence, if all individuals in the population expect a $z'=0.2$ fraction of the population to adopt perpetration, then that fraction will indeed adopt and individuals with number names above $z'=0.2$ will not. Hence, the shared expectation $z'=0.2$ will be fulfilled. But note that there is another self-fulfilling z equilibrium in Figure 4, namely at $z''=0.8$. If all individuals in the population expect that fraction of adoption, then all individuals with number names between 0 and 0.8 will adopt. The last adopter in the segment will be person 0.8, whose reservation valuation just meets the cost $r(z)f(z)=(1-z)(12.5z)=(1-0.8)(12.5)(0.8)=2=c$. Persons with names above 0.8 will not adopt perpetration because their reservation valuations will be less than 2. For example, person 0.9’s reservation value will be $(1-0.9)(12.5)(0.9)=1.125<c=2$. Hence, Figure 4 shows three possible self-fulfilling expectations equilibria, namely at $z=0$, at $z'=0.2$, and at $z''=0.8$.

[Figure 4 about here]

Continuing with Figure 4, explore now the stability and tipping point properties of the three equilibria. Consider the equilibrium given by the origin point in Figure 4. This is a *stable*

equilibrium. Suppose the population temporarily adopts a small positive expectations fraction $z = \varepsilon > 0$. In this case, the highest value of atrocity perpetration on the curve would be less than the cost, so the highest valuing individual, and thus all others, would not participate. In this case, the outcome would move back to the origin. Consider next the $z' = 0.2$ equilibrium. Suppose the population temporarily adopts a slightly smaller expected fraction of atrocity perpetration $z = z' - \varepsilon < 0.2$. Since the curve in Figure 4 lies below the cost line at $z' - \varepsilon < 0.2$, the highest valuing individual would not adopt atrocity perpetration and, thus, nobody else would either. As such, the outcome would gravitate to the origin. Thus, $z' = 0.2$ is an *unstable equilibrium*. But suppose the population temporarily adopted a slightly *larger* expectations fraction $z = z' + \varepsilon > 0.2$. Since the curve in Figure 4 lies above the cost line at $z' + \varepsilon > 0.2$, individual $z' + \varepsilon$ prefers adoption over abstention and thus becomes an adopter. Since that individual has come on board, network benefits from adoption have been notched up ($f' > 0$), so that the next individual above $z' + \varepsilon$ will adopt and so on until the last adopter is reached in chain reaction fashion, namely person z'' . Yet individuals with names above z'' would have reservation valuations from adoption that are less than the cost of adoption as shown by the portion of the curve on the right side of Figure 4 lying below the cost line. In sum, these results establish that z' is an unstable equilibrium and $z = 0$ and z'' are stable equilibria.

Four important points follow in regard to atrocity propagation and prevention in Figure 4. First, equilibrium z' (being unstable) is a *critical point* or a *tipping point* for the “success” of atrocity propagation. If the number of atrocity adopters somehow reaches $z' + \varepsilon$, then upward pressure (or demand) for additional agents to become perpetrators will be set into motion as the society moves from the $z' + \varepsilon$ fraction of adoption to a much higher rate of participation at z'' . Second, atrocity architects will attempt to get their societies “over the hump” of resistance to

atrocities such that the much higher equilibrium of participation occurs. Third, by manipulating the elements of the model, atrocity architects can attempt to achieve the high participation they seek. For example, by lowering the cost of atrocity participation (reduce c), the z' equilibrium will decline (move left in the figure), thus making it easier to get over the hump. Moreover, the z'' equilibrium will increase (move right), thus implying a greater rate of participation once the hump is gotten over. If cost is unchanging or given, atrocity architects nonetheless could attempt to shift the curve in Figure 4 upward by increasing actors' intrinsic valuations for atrocity, r , and by promoting greater network benefits from atrocity, f . Just as if costs had been reduced, such actions also would serve to reduce z' and increase z'' . Finally, fourth, atrocity preventers want to do the opposite of the architects; specifically, they would like the cost, c , to be high enough, and the r and f valuation functions to be low enough, so that the cost line in Figure 4 lies everywhere above the curve.

2.3 Derivation of S-Shaped Diffusion Curve

Following the groundwork that has been laid, we can now derive an S-shaped diffusion curve based upon the foregoing modeling. Recall that z symbolizes the expected fraction of the population that will adopt atrocity perpetration and that this expectation is shared by all individuals from the in-group. Just because z is what everybody expects does not imply that z will be the *actual* rate of adoption at a point in time. Let \hat{z} be the fraction of the population that solves the equation $r(\hat{z})f(z) = c$ where all other terms are as described above. Note that since person \hat{z} just adopts atrocity (because \hat{z} 's reservation valuation is just equal to the cost), all individuals from 0 up to \hat{z} will adopt also. Assuming the functional forms for r and f from above, the equality condition then is:

$$(1 - \hat{z})(12.5z) = c. \tag{2}$$

Solving (2) for \hat{z} gives:

$$\hat{z} = 1 - \frac{c}{12.5z} . \quad (3)$$

Note in (3) that if the term to the right of the minus sign is greater than 1, then \hat{z} is negative and falls outside the bounds of the model. Hence, equation (3) governs the relationship between the shared *expectation* of the population, z , and the *actual* rate of adoption, \hat{z} , assuming $c/12.5z \leq 1$. If this condition does not hold, then $\hat{z}=0$. More generally, let the right-hand side of the equality in (3) be represented by the $g(z)$ function. The general relationship between \hat{z} and z will then be:

$$\begin{aligned} \hat{z} &= g(z) & , \text{ when the restraint condition holds} \\ \hat{z} &= 0 & , \text{ otherwise.} \end{aligned} \quad (4)$$

We demonstrate how the model now leads to an S-shaped curve in Figure 5 based upon equation (3) and the restraint condition $c/12.5z \leq 1$ under the assumption that $c=2$. First, note that the restraint condition will not be met for values of $z < 0.16$. Hence, for population-wide expectations of $0 \leq z < 0.16$, the actual rate of adoption will be $\hat{z}=0$. We are then led to the diffusion curve shown in Figure 5 (the 45^0 line will be explained shortly). On the X axis, we measure at a point in time the population's shared expectation of atrocity adoption, z . The diffusion curve then shows, given z , what the actual outcome of atrocity adoption, \hat{z} , will be as measured on the Y axis. For example, if $z=0.1$ as shown on the graph, then the actual atrocity adoption, \hat{z} , in that period will be $\hat{z} = 0$. Suppose $z=0.3$ as shown on the graph. The diffusion curve then shows that the actual atrocity adoption rate will be $\hat{z} > 0.3$ because the diffusion curve lies above the 45^0 line when $z=0.3$. Assuming myopic adjustment in expectations among the population (not an unreasonable assumption for a large population), the new expectation of atrocity adoption in the next period will be $z=\hat{z}$ as traced over to the 45^0 line (call that z -value z_2). As the graph shows, when $z=z_2$, the actual atrocity adoption rate will be $\hat{z} > z_2$ because the

diffusion curve lies above the 45^0 line when $z=z_2$. This process of expectations and valuation adjustment will continue until the expected and actual rates of adoption are equal at $\hat{z} = z'' = 0.8$. Figure 5 implies stable equilibria at z values of 0 and 0.8 and a critical point (unstable equilibrium) at a z value of 0.2.

[Figure 5 about here]

2.5 Comparative Statics

As before, the comparative static properties of the diffusion curve in Figure 5 boil down to elements that would change the cost of atrocity adoption, c , shift or rotate the intrinsic valuation curve, r , or shift the network valuation curve, f . For example, suppose the cost of atrocity adoption increases from $c=2.0$ to $c=4.0$, all else equal. This will cause the diffusion curve to lie everywhere below the 45^0 line as shown in Figure 6. The result will be no acceptance of atrocity in the population at all, i.e., there will be a unique and stable equilibrium at $\hat{z} = z = 0$. The exact same diffusion curve and result would obtain if c remained at 2 but the intrinsic valuation for adoption function was cut in half, from $r=1-x$ to $r=0.5(1-x)$. If instead, the c value and the r function were unchanged but the value for adoption from network effects was cut in half, from $f=12.5z$ to $f=6.25z$, again the exact same diffusion curve and result would obtain as in the figure. Of course, if c was to fall enough and, simultaneously, the valuation functions increase enough, then a result could emerge in which the diffusion curve lies mostly or even everywhere above the 45^0 line, thus leading to a high or complete diffusion of atrocity adoption. For example, in Figure 7 the cost of atrocity adoption is reduced from $c=2$ to $c=1$ and the networking function changed to $f(z)=3+12.5z^2$. Reflecting perhaps a not unrealistic scenario, note that for this new networking function, networking benefits will exist even when $z=0$ (i.e., $f(0)=3$) and networking benefits increase at an increasing rather than constant rate as given by the z^2 term. The g function

becomes $g(z)=1-(1/12.5z^2)$. The result is an extremely high fraction of adoption at $z^*=0.93$. This outcome is a unique and stable equilibrium.¹

[Figure 6 about here]

[Figure 7 about here]

2.6 Statement of the General Model

We can now state a general formulation of the S-shaped model. The diffusion curve can come in many shapes, but such curves will be generated from the spontaneous or intrinsic valuation function of individuals in the population, i.e., the $r(x)$ function in our case, *and* the networking benefits function, $f(z)$ in our case. The aggregated benefits function could be multiplicative, i.e., $r(x)f(z)$, but this need not be so. Many functional forms for the aggregate benefits function are possible depending on the functional forms for r and f . Assuming reasonable properties of the r and f functions (specifically $r'<0$ and $f'>0$),² a diffusion curve will be implied that will generally be nonlinear, will have an S-shape (i.e., with segments that at first increase at an increasing rate and then increase at a diminishing rate), and will lie everywhere above, everywhere below, or intersect the 45° line one or more times. When the diffusion curve intersects the 45° line from below, the result will be an unstable equilibrium, which will be a tipping point on either side of which the population behavior moves to a stable equilibrium. When the diffusion curve intersects the 45° line from above, the result will be a stable equilibrium as deviations from that point in either direction move the dynamics of the behavior back to the equilibrium. The dynamic behavior of atrocity adoption in the population can be analyzed graphically as we have done

¹ Unlike an earlier example, given the specific functional forms that we have been working with, 100 percent adoption will not occur unless the cost of adoption is zero. Other functional forms, however, could give rise to the diffusion curve lying everywhere above the 45° line even at positive cost, thus giving a unique and stable equilibrium at $z^*=1$.

² The expression $r'<0$ refers to a negative (falling) slope and means that the lower the cost, c , the higher the rate of adoption. The expression $f'>0$ refers to a positive (rising) slope and means a rising acceptance of atrocity, given that peers have accepted.

here. Although pictorial, such analysis is “nevertheless completely rigorous” (Easley and Kleinberg, 2010, p. 461). Many pictorial examples of S-shaped diffusion curves in the literature, for example those of Schelling (1971, 1978) and Granovetter (1978), can be conceived of as examples that fit the framework laid out here.

3. Models of Diffusion across Networks

Both the Bass and the generalized S-curve models analyze diffusion with a *population*, and thus ignore detailed structural aspects of social relations among networked individuals. Here we explicitly introduce *network structure* into models of diffusion. Our purpose is to apply network models of diffusion to mass atrocity acceptance and mass atrocity rejection across networks (or neighborhoods) of individuals.

3.1 Basic Cascade Model of Atrocity Acceptance and Rejection across a Network

Consider the coordination game between two players (1 and 2) in Figure 8. The two actions available to each player in the game in our context are aggression (A) and peace (P) directed against an out-group. If agents 1 and 2 each direct aggressive actions against an out-group, they each receive a positive payoff $a > 0$. Likewise, if each directs peaceful actions against an out-group, each receives a positive payoff $b > 0$. However, if the two agents choose opposite actions (and thus are uncoordinated), each receives a payoff of 0.

[Figure 8 about here]

Again following Easley and Kleinberg (2010, p. 500), assume in Figure 9 that player 1 would like to maximize its payoff from the pairwise interactions with its five neighbors (players 2–6). Two of player 1’s neighbors (2 and 3) have chosen to be aggressive toward the out-group (they have chosen “A” in Figure 8). Player 1’s other three neighbors (4, 5, and 6) have chosen to be peaceful toward the out-group (they have chosen “P” in Figure 8). Assuming that player 1

must choose A or P and cannot choose any kind of mixture of the two, what choice would player 1 make to maximize its payoff? Easley and Kleinberg (2010, p. 501) demonstrate that if at least the fraction $f=b/(a+b)$ of player 1's neighbors choose A, then player 1 will maximize its payoff by choosing A also. On the flip side, if less than the fraction f of 1's neighbors choose A, then 1 maximizes its payoff by choosing P. As a numerical example, suppose that $a=4$, $b=2$. In this case, the actual fraction of 1's neighbors choosing A in Figure 9 is $2/5=0.4$ and thus greater than $f=2/6=0.33$. Hence, player 1 will maximize its payoff by choosing A. If, instead, $a=2$, $b=2$, the actual fraction of 1's neighbors choosing A (0.4) would be less than $f=2/4=0.5$, so A would maximize its payoff by choosing P. Crucially, Figures 8 and 9 demonstrate that payoffs from social interactions *and* network structure affect what player 1 will do to maximize its payoff. These principles can now be generalized to model the contagion of mass atrocity acceptance across varying network structures.

[Figure 9 about here]

Figure 10 shows friendship ties in a neighborhood of 12 individuals. Assume a crisis hits the country in which this neighborhood resides and an authority group attempts to initiate atrocity against an out-group. (Neighbor and neighborhood are general terms that need not imply a spatial or location neighborhood but can refer to a social network of family, friends, coworkers, and so on.) All of the individuals in Figure 10 are from the in-group. In panel (a) we assume initially that each individual adopts a posture of peace (P) toward the out-group (ignore the dark circles in panel (a) for the moment). Based on Figure 8, this implies that each peaceful neighbor achieves a payoff of b from each of its tied neighbors (e.g., player 1 chooses P and has two linked neighbors who choose P, so player 1 would receive a payoff of $2b$). Suppose now that the a and b parameters in Figure 8 are such that the key fraction determining when an individual's

choice of aggression (A) is optimal is $f=0.5$. This is a relatively malignant threshold in the sense that an individual will find it optimal to choose A when only half (or more) of its neighbors choose A. Perhaps government authorities have instituted rewards and punishments in this society such that a is relatively high and b relatively low in Figure 8, thus giving rise to $f=0.5$.

Assume now in panel (a) of Figure 10 that, for exogenous reasons, individuals 5 and 10 spontaneously switch to aggression (A). This is indicated by the darkened circles for these agents in the figure. Perhaps individuals 5 and 10 were swayed by government propaganda or they were given an exogenous side payment to switch to A. Whatever the case, a contagion of atrocity acceptance will now unfold. Since 5 is darkened, now half of the neighbors of individuals 3, 4, 6, and 7 have chosen A, and so they, too, will choose A (given $f=0.5$). In panel (b) we thus darken the circles for agents 3, 4, 6, and 7 to indicate their conversions to atrocity acceptance. It now follows that individuals 1, 2 and 8 have at least half of their neighbors choosing A, so in panel (c) they choose A also. In the next round, individual 9 will have at least half of its neighbors as atrocity accepters, so 9 will choose A. Finally, the last set of individuals (11 and 12) will convert to A for like reason. Hence, we arrive at panel (d) in which all circles are darkened and the whole network has become infected with atrocity acceptance.

[Figure 10 about here]

The model described in Figure 10 is a complete model in that it specifies *initial conditions* (agents 5 and 10 adopt A, all others adopt P), a *threshold rule* rooted in a stage game ($f=b/(a+b)$, which was set equal to 0.5), and the *dynamic* unfolding of contagion until the process ends (Easley and Kleinberg, 2010, pp. 501–2). Of course, the specific network structure in Figure 10, the stage game in Figure 8, and the threshold rule could be quite different and even individualized (i.e., each agent could have its own threshold rule), but the process of working

through the contagion in the model would follow the same procedure used here. Given that we have a complete model in this sense, we can move on to comparative statics on the model.

3.2 Comparative Statics of the Basic Cascade Model

In what follows we consider three comparative statics: (1) change in initial conditions, (2) change in payoffs in the stage game and thus change in the threshold rule, and (3) change in network links among existing agents. Later we work in additional comparative statics; in particular, we will consider immunized agents and key player policy when we use the model in Figure 10 to analyze atrocity prevention. Other types of comparative statics—e.g., one-way or directed links, incomplete information on links, multiplex links (i.e., some agents choosing A vis-à-vis some neighbors and P vis-à-vis other neighbors)—are left for future research.

Return to Figure 10(a), but with the proviso that we change one of the initial (exogenous) adopters of aggression. Suppose now that agent 4 (along with agent 10) is an initial adopter of A rather than agent 5, all else equal. This scenario is shown in Figure 11. Now run the contagion process forward through time. Note that none of the agents in the upper cluster of players surrounding agent 4 have at least half of their networked neighbors as adopters of A. This is also the case for the middle and lower clusters of agents in the figure. In short, there is no contagion in Figure 11! This simple result is remarkable. In moving from panel (a) in Figure 10 to Figure 11, we have not changed the number of the initial adopters of A at time zero, the total number of individuals in the network, the structure of neighbor links in the network, or the threshold rule for the adoption of A. We merely changed the location on the network of one of the initial adopters. Despite this truly trivial change, the effect on the aggregate outcome is dramatic. Specifically, this small change in an initial condition caused the aggregate outcome to switch from complete contagion to no contagion at all. This captures puzzles 1 to 7 but now in a more

complexly specified network and diffusion model. Recall that puzzle 1 asked about how small-scale atrocity adoption can scale up to mass participation. The details for the diffusion process address this. Puzzle 2 asked why one location experiences mass participation while another does not. Differences in initial network conditions across different locations can explain this. Puzzle 3 asked why individual participation may “flip” from adoption to nonadoption. A change in conditions, such as when an initially adopting agent moves from location 5 to location 4, can explain this. Puzzle 4 asked why individuals can be both persecutors *and* rescuers. What can explain this is that agents are multiplex, that is, they are involved in a variety of networks, each with its own conditions, threshold rules, and specific network structures. Puzzle 5 asked why some people, who are strongly opposed to acts of mass atrocity, nonetheless help commit them. Network theory suggests that peer effects, $f(z)$ or pairwise interaction payoffs, in the network may be strong enough to overcome an individual’s reservation value, $r(x)$. Puzzle 6 asked why, post-atrocity, some individuals appear to exhibit genuine shock and remorse at their own participation. Again, a change in network structure, such as when key agents are removed from the network structure, say in the wake of Germany’s defeat in the Second World War, can explain this. Puzzle 7 asked what explains the “silence of the majority” in many mass atrocity cases, and network theory suggests that particular combinations of initial conditions, threshold rules, and network structure can generate such outcomes. The “reductionist” point of models such as those described here is to learn whether, beyond case-specific studies of mass atrocities, there are *general* factors common to all (or most) mass atrocities on which general policy approaches may focus, at least as a first cut at a specific mass atrocity situation.

[Figure 11 about here]

Next, in Figure 12 we take as our initial starting point at time zero the situation depicted in Figure 11 in which agents 4 and 10 exogenously adopt behavior A, but we change the threshold rule from $f=1/2$ to $f=1/3$. Panel (a) of Figure 12 is a reproduction of Figure 11. With the even more malignant threshold rule now operative (only $1/3$ instead of $1/2$ of neighbors need to adopt before the agent in question adopts, perhaps owing to changes in payoffs in the stage game in Figure 8 by the atrocity architects), note in panel (b) that agents 1 and 12 now have at least one-third of their neighbors as atrocity adopters, so they choose A also. Hence, we darken the circles for agents 1 and 12 in panel (b). It then follows in panel (c) that agents 2, 5, 9, and 11 have at least one-third of their neighbors as adopters, so they adopt too and their circles become darkened. Now in panel (d), all of the remaining agents (3, 6, 7, and 8) have at least one-third of their neighbors adopting, so we darken their circles in panel (d). The result is one of complete adoption of atrocity acceptance on the network owing to a change in the threshold rule, all else constant.

[Figure 12 about here]

In Figure 13 we return to our initial setup in which agents 5 and 10 are initial (exogenous) adopters and the threshold rule is $f=1/2$. Recall from Figure 10 that these conditions led to the complete contagion of atrocity acceptance on the network. Now we add to that scenario what again seems like a trivial adjustment. Specifically, we add a new neighborhood tie between agents 3 and 4 in Figure 13 as shown by the dashed line between them. Now trace the contagion on the network for this new scenario. In panel (b), agents 6 and 7 have at least half of their neighbors as adopters, so we darken their circles as they convert to adoption. Now agent 8 has at least half of its neighbors adopting, so it adopts too leading to a darkened circle for 8 in panel (c). Finally, in panel (d) agent 9 converts which then facilitates the conversions of 11 and

12 to adoption. But note that none of the individuals in the top cluster of panel (d) convert. This is a surprising result, namely, that just one additional neighborhood tie among peace types in the top cluster allowed the whole cluster to be buffered against atrocity acceptance. In this example, a slightly denser network among peaceful types enables the cluster to resist being infected by the atrocity acceptance contagion.³ We will now see how this result can be generalized to any network structure with undirected ties and a threshold rule.

[Figure 13 about here]

3.3 Clusters as Obstacles to Cascades: A General Statement

The previous subsection shows how the degree of contagion across a specifically contrived network can vary dramatically depending on seemingly small, even trivial, changes (“trembles”) in initial conditions, threshold rules, network structure, or number of links. Whether extensive or limited contagion of atrocity acceptance or rejection occurred on the network seemed almost idiosyncratic in the sense that a slight twist in a condition here or there could drastically alter the diffusion process. A question naturally arises: Can a more general statement be made about the diffusion (or lack thereof) of a substance across *any* network? More generally, we will show that groups within networks characterized by *homophily* (a relatively close or tightly-knit community) serve as “barriers to entry” to new behavior. Already we have seen an example of this idea in Figure 13. Specifically, when we added one neighborhood (or friendship) link, namely between agents 3 and 4, to the top cluster of individuals in the figure, we made that

³ The contagion of atrocity acceptance across locales during mass atrocities is often dramatically uneven. Some locales seem immune to atrocity acceptance while others give way to a high level of atrocity actions quite rapidly. For the 1994 Rwandan genocide, McDoom (2014) presents empirical evidence that “cohesive communities resist[ed] elite attempts to divide them” (p. 34). Similarly, McDoom’s (2013) empirical evidence implies that as the number of peaceful to violent people in an individual’s neighborhood or household increases, the likelihood of this individual’s participation decreases. Network theory provides plausible explanations for such observations.

group within the network a slightly more tightly-knit community. In the example, the slight increase in the group's tightness caused it to become a barrier to atrocity acceptance.

Following Easley and Kleinberg (2010, p. 502), we conceptualize a group of tied individuals within a network as a cluster of individuals with a certain degree of tightness or interconnection among them. Specifically, a cluster with density d is a set of actors such that each actor in the set has at least the ratio r of its tied friends in the set. Consider, for example, Figure 14 which is a reproduction of panel (a) of Figure 13 but with the tie between agents 3 and 4 now a solid rather than a dashed line. Recall that the idea of atrocity acceptance was not able to break into the top cluster of agents (1–4) (see panel (d) of Figure 13). The density of that top cluster in Figure 13 is $d=2/3$, that is, each agent 1–4 has at least two-thirds of its friendship ties within the cluster of agents 1–4. Let agents 5–8 be a cluster as well. The density of that cluster is $d=1/2$. Finally, let individuals 9–12 be a cluster, with density of $d=1/3$. Hence, the cluster that was most dense in Figure 13 was the one cluster in which the contagion of atrocity acceptance did not spread given the initial conditions and the threshold fraction f . Meanwhile, the other two clusters in Figure 13 were not sufficiently dense to prevent the contagion given the initial conditions and the threshold fraction f .

[Figure 14 about here]

In contexts such as the models depicted in this section, Easley and Kleinberg (2010, p. 507) make the following claim:

Claim: Consider a network in which there is a set of initial adopters of behavior A (atrocity acceptance in our context) with a threshold of fraction f for adopting A for remaining agents in the network.

(i) If the remaining network contains a cluster of density greater than $1-f$, then the set of initial adopters will not cause a complete cascade.

And

(ii) For threshold f , if the initial set of adopters does not cause a complete cascade, the remaining network contains a cluster of density greater than $1-f$.

(For a proof, see Easley and Kleinberg, 2010, pp. 507–9.)

Note how the Easley/Kleinberg (E/K) statement is fulfilled in Figure 14. With $f=1/2$, the network does contain a cluster with density greater than $1-f=1/2$, namely, the top cluster with density $d=2/3$. It then follows that both (i) and (ii) in the E/K claim are satisfied. Note also how the E/K claim is supported in Figures 10 and 12 above. In both figures, the density of the top cluster is $d=1/2$ rather than $d=2/3$. Part (i) of the E/K claim does not hold for those figures, and a complete cascade occurs for each.

The more general principle associated with the E/K model is that tightly-knit communities serve as barriers to entry to contagion. This can be good news or bad news for atrocity contagion. Just as a tight-knit community might be more likely to resist atrocity acceptance, a tight-knit community could also be more resistant to peacefulness toward an out-group. As Amartya Sen (2006, p. 1) writes: “identity can also kill—and kill with abandon.”

4. Network Tentacles of Atrocity Perpetration

The models covered up to this point focus on the diffusion of atrocity *acceptance* or the resistance to atrocity across a population or within a neighborhood (locale) of individuals. The people over which such diffusion processes play out are assumed to be from the in-group. Even within a population and even within neighborhoods in which atrocity becomes acceptable, there is still the matter (from the architect’s point of view) of *bringing force to bear to destroy* the out-

group. The bringing of such force involves networks of individuals and organizations including troops to carry out the killing and prevent the victims from fleeing, bureaus and logisticians to manage equipment and other supplies for the troops, commanders to organize operations, and businesses to provide the implements of people-group destruction. Much like the “tooth-to-tail ratio” is used in the military planning literature to summarize the many facets of networks involved in preparing for and conducting war, we conceive of “networked tentacles of atrocity perpetration” along similar lines. In this section, we offer a stylized linear quadratic model of network tentacles of atrocity perpetration.

4.1 Basic Setup of the Linear Quadratic Model of Atrocity Tentacles

Figure 15 is a highly stylized depiction of a social network designed to bring tentacles of out-group destruction to three villages, A, B, and C. At the center of the network is player 1 who is the atrocity architect. Moving out from the center are three people, players 2, 3, and 4, who are the regional managers or bureau heads associated with villages A, B, and C, respectively. Finally, agents 5–10 represent commanders of troops carrying out atrocities in their assigned locales. For the moment, ignore the w numbers in the figure.

Assume that each individual player, i , in the network in Figure 15 chooses a level or intensity of an action, $x_i \geq 0$, assumed to be action or harm against the out-group. Following the linear quadratic model (LQM) from the social networking literature (e.g., Jackson, 2008), each player’s utility or payoff is given by:

$$U_i = a_i x_i - \frac{b_i}{2} x_i^2 + \sum_{j \neq i} w_{ij} x_i x_j , \quad (5)$$

where $a_i \geq 0$ and $b_i > 0$ are benefit and cost scalars, respectively, and $w_{ij} \geq 0$ is the weight or importance that player i places on player j ’s action (Jackson, 2008, p. 290). Equation (5) implies that each unit of x_i brings to player i marginal benefits of $a_i + \sum_{j \neq i} w_{ij} x_j$ (i.e., additional benefits

per additional unit of action) and marginal costs of $b_i x_i$. The w_{ij} parameter captures strategic complementarities (when $w_{ij} > 0$) among linked agents. This aspect of equation (5) captures the increase in i 's self-perceived well-being when j 's positive action interacts with i 's positive action ($x_i x_j$). The reasons for i 's increased well-being from j 's action could be multiple, including feelings of comradery from having a person that one is directly connected to operating for the “cause” as well—that is, a peer effect in which my linked counterpart’s higher action causes me to want to engage in a higher action too (e.g., to show myself well within the in-group)—or information flows among networked agents that enhance the “ideological necessity” of destroying the out-group.

Assume now that each agent i maximizes U_i in (5) by choosing x_i , with all other elements in (5) treated parametrically. This leads to the following reaction function for i (Jackson, 2008, p. 291):

$$x_i = \frac{a_i}{b_i} + \sum_{j \neq i} \frac{w_{ij}}{b_i} x_j \quad \text{or} \quad x_i - \sum_{j \neq i} \frac{w_{ij}}{b_i} x_j = \frac{a_i}{b_i}. \quad (6)$$

The system of reaction functions then can be written in matrix algebra format as:

$$Ax = B \quad (7)$$

where A is the 10×10 matrix of coefficients that multiply the x variables, x is the 10×1 vector of x_i variables, and B is the 10×1 vector of a_i/b_i terms. The solution to (7) is

$$x^* = A^{-1}B, \quad (8)$$

where A^{-1} is invertible and the inverse of A and x is nonnegative.

4.2 Numerical Example of LQM Model

To illustrate the workings of the model in the previous subsection, we begin with the assumption that the odd-numbered commanders (5, 7, and 9) favor atrocities against the out-group but the even-numbered commanders (6, 8, and 10) are uncomfortable bystanders: They would rather not

carry out atrocities but also do not want to disobey orders and try to rescue victims. To give specific numbers to Figure 15 and to the equations of the model, assume that $a_i=2$, $b_i=1$, $w_{ij}=0.2$, and $w_{ii}=0$ for $i=1-5, 7$, and 9 and $a_i=0$, $b_i=1$, $w_{ij}=0$, and $w_{ii}=0$ for $i=6, 8$, and 10 . The numerical values for the w terms are labeled in Figure 15. Equation (7) would then be:

$$\begin{bmatrix} 1 & -0.2 & -0.2 & -0.2 & 0 & 0 & 0 & 0 & 0 & 0 \\ -0.2 & 1 & -0.2 & -0.2 & -0.2 & 0 & 0 & 0 & 0 & 0 \\ -0.2 & -0.2 & 1 & -0.2 & 0 & 0 & -0.2 & 0 & 0 & 0 \\ -0.2 & -0.2 & -0.2 & 1 & 0 & 0 & 0 & 0 & -0.2 & 0 \\ 0 & -0.2 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & -0.2 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & -0.2 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 \\ x_8 \\ x_9 \\ x_{10} \end{bmatrix} = \begin{bmatrix} 2 \\ 2 \\ 2 \\ 2 \\ 2 \\ 0 \\ 2 \\ 0 \\ 2 \\ 0 \end{bmatrix}. \quad (9)$$

Following equation (8), the solution to (9) is:

$$\begin{bmatrix} x_1^* \\ x_2^* \\ x_3^* \\ x_4^* \\ x_5^* \\ x_6^* \\ x_7^* \\ x_8^* \\ x_9^* \\ x_{10}^* \end{bmatrix} = \begin{bmatrix} 5.81 \\ 6.36 \\ 6.36 \\ 6.36 \\ 3.28 \\ 0 \\ 3.28 \\ 0 \\ 3.28 \\ 0 \end{bmatrix}. \quad (10)$$

The aggregate amount of actions in (10) is $X^*=36.7$. (We note that the interpretation of the amount of harm caused is unclear. A simple assumption, given symmetry in Figure 10, is to divide 36.7 by 3 (=12.2) and assume that this is the amount of harm brought to each village.

Another interpretation is to assume that the actions of the architect, player 1, is purely public such that $x_1^*=5.81$ applies to each of the three tentacles. This then would give $[(36.7-5.81)/3] + 5.81 = 16.1$ units of harm directed to each village. Although obviously relevant in the real world, the interpretation does not matter for our modeling purposes here.)

The atrocity network in Figure 15 and the solution in equation (10) assume that the even-numbered commanders are not supportive of atrocity (their a_i values were zero and their link weight terms with others were zero). Assume now that such commanders are replaced with “willing executioners” (Goldhagen, 1996). For simplicity, we assume all commanders now have the parameter values of the odd-numbered commanders in equation (9). The aggregate output of the tentacle model rises from $X^*=36.7$ units of harm to $X^*=62.6$ (a 71 percent increase), a rather large numeric increase on network output from a relatively small policy change of replacing three unwilling commanders in the field.

4.3 Comparative Statics Analysis of LQM Model

We now summarize various comparative static results for the solution of the general LQM model in equation (8) and also with numerical examples. Following Jackson (2008, pp. 292–3) and Ballester, Calvó-Armengol, and Zenou (2006), assume $a_i=a$ and $b_i=b$ for all agents i , but continue to assume that the weight parameters $w_{ij} \geq 0$ can be heterogeneous. Under these conditions, it can be shown that the equilibrium levels of harmful actions in (8) can also be represented by (see Jackson, 2008, p. 292):

$$x^* = \frac{a}{b} \left(I - \frac{1}{b} w \right)^{-1} \frac{a}{b} \mathbf{1}, \quad (11)$$

where I is the identity matrix, w is the weight parameters matrix, and $\mathbf{1}$ is an $n \times 1$ vector of ones.

From (11), an increase in the benefit parameter a or a decrease in the cost parameter b will increase each agent's x value in the vector x^* . It also follows that an increase in any one of the w_{ij}

terms in matrix w will increase the equilibrium level of harm for each agent in the network who has a directed (one-way) path to agent i (see Jackson, 2008, p. 292 for a proof).

Condition (11) is also amenable to the analysis of Key Player Policy (KPP) (Ballester, Calvó-Armengol, and Zenou, 2006). The key player is the one whose removal causes the production on the network to decrease the most assuming that the remaining players re-optimize their actions. Ballester, Calvó-Armengol, and Zenou (2006) identify the key player in a condition like (11) based upon a Bonacich (1987) network centrality measure, C^B . The Bonacich centrality measure for an actor counts the number of all of the paths that emanate from that actor's node, weighted by a decay factor so that an agent's reach decreases with the lengths of the paths (Ballester, Calvó-Armengol, and Zenou, 2006, p. 1404). Following Ballester, Calvó-Armengol, and Zenou (2006) and Jackson (2008, p. 292), assume that the a and b terms in the LQM are the same for each actor and that any heterogeneity on the network comes only through the link weights, w_{ij} . Further assume that $1/b$ is small enough so that C^B is well-defined and nonnegative (Jackson, 2008, p. 292). It follows that the equilibrium levels of harm on the network are (Jackson, 2008, p. 292):

$$x^* = \frac{a}{b} (\mathbf{1} + C^B), \quad (12)$$

where $\mathbf{1}$ is an $n \times 1$ vector of ones, n is the number of actors on the network, and C^B is a vector of Bonacich centrality measures, which in turn are functions of the b cost parameter and the matrix of weight links, w . Ballester, Calvó-Armengol, and Zenou (2006) show that the largest reduction in the network's output occurs from the removal of the actor with the highest value of a variant of the Bonacich centrality measure (see Jackson, 2008, pp. 292-293). While we do not pursue these technical matters here (as indeed the literature knows of multiple centrality measures, each with its own good uses), the point is that network theory can help identify key players whose

removal from the network will damage its productivity. For a high-level empirical application of Key Player Policy to fighting among armed groups in the Democratic Republic of the Congo (DRC), and its attendant atrocities, see König, *et al.* (2017).

5. Contesting Networks of Atrocity Perpetration and Prevention

There are four general types of interventions that third parties could attempt to impose on a network to prevent atrocity: (1) Deployment of various “carrots and sticks” to change the benefits and costs of individuals on the network such that less hostility is generated (i.e., atrocity preventing comparative statics), (2) nullification of one or more key players on the network, and/or other centrally important actors, and/or some or all of the links among atrocity supporting actors, (3) insertion of resistance actors (saboteurs) at key places in the network, and (4) insertion of a third party’s “tentacles of atrocity prevention” to directly contest the atrocity-producing organization’s work. In this section, we model the latter of these alternatives as well as providing brief discussion of (1) and (2).

Assume an atrocity perpetrating network has $i=1, \dots, n$ atrocity tentacles in which $x_i \geq 0$ units of harm are being directed per tentacle against an out-group. Think of the i ’s as locations in geographic space and/or in time. Similarly, assume a network of third party helpers directs $y_i \geq 0$ units of atrocity resistance to each of the perpetrator’s tentacles to directly contest the efforts of the perpetrators. Assume the result of any contestation between the atrocity perpetrators and the atrocity preventers is governed by the following Tullock-like contest success function (CSF):

$$r_i^S = \frac{y_i}{x_i + \tau_i y_i}, \quad (13)$$

where r_i^S is the ratio of the vulnerable population at location i that is saved from victimhood and $\tau_i \geq 0$ is a technology of contestation parameter appropriate to location i , which depends on

geography and other elements that aid or hinder the protection of vulnerable people. Consistent with CSFs deployed in the literature, $r_i^S(x_i, y_i) = r_i^S(0,0) = 1$.

As a numerical example of the contestation model, return to the atrocity perpetrator's network in Figure 15 under the assumption that all ten actors are perpetrators. Recall that the network's aggregate number of units of harm directed to the out-group was $X^*=62.6$. Assume each of the three tentacles of atrocity receive one-third of those units, or about 20.9 units of harm for each tentacle, which we will round to 21 units for simplicity. Suppose a network of third party actors generates a total of $Y^*=42$ units of action to directly contest the 63 units of harm of the atrocity network. If the third party deploys 14 units of contestation to each atrocity tentacle and if $\tau_i=1$ for each i , then 40 percent of the vulnerable population at each tentacle location will be protected and 60 percent will be lost.⁴

A second approach to modelling contesting networks is to assume that each network produces an aggregate output. Such outputs can be thought of as efforts or "weapons" in the contest between the networks in which the outcome of the struggle would be modeled by a contest success function (CSF). For example, the aggregate output of the atrocity-perpetrating network is $X^*=62.6$ when all commanders are on board. Suppose a third party brought $Y=50$ units to contest the actions of the atrocity-perpetrating network. A simple ratio-form CSF would be $p=Y/(X+Y)=50/(62.6+50)=0.44$ where p is the proportion of the vulnerable civilian population that survives the atrocity network's assault or the probability that any given vulnerable civilian survives. In this example, 44 percent of the vulnerable civilian population would survive, but 0 percent would survive if there was no third party help ($Y=0$). Once more, we do not pursue these

⁴ An alternative interpretation is that there will be a 0.4 probability of survival for each vulnerable person at each location and the expected values will be that 40 percent of each vulnerable population survives at each location.

matters here and only point to the potential value of bringing explicit network theory-induced reasoning to bear on real world cases of mass atrocities such as genocides.

6. Conclusions

We began the article with seven puzzles and, in the text that followed, showed that concepts from network theory help address all seven puzzles. Even though each real-world case of mass atrocity is somehow “special” (hence the very many case histories), all of them would seem to emerge from a unified underlying structure, thus linking highly detailed “micro” level studies to a “macro” level *general theory of mass atrocity*. The theory identifies key elements regarding initial conditions, diffusion processes, adoption and imitation rates, threshold and neighborhood effects, network structure, multiplex networks, cost and benefit considerations (including psychological costs and benefits such as feelings of animosity or amity toward others), and so on. Depending on the particular constellation of these and other parameters in one or more networks, the promise of network theory is that real-world cases of mass atrocity can be “recreated” theoretically. To the degree that this effort succeeds, it is then possible to also link what is known from case studies and other efforts about mass atrocity intervention and prevention to network theory, as we have indicated in several instances above (e.g., the Key Player Policy or location-specific interventions on a network).

This article has only begun to apply network concepts and models to mass atrocity onset, spread, and prevention. Numerous extensions are possible. For diffusion over a population (and prevention of diffusion), much research terrain is still to be explored on adapting *SI*, *SIR*, *SIS*, and *SIRS* models from the epidemiology literature to mass atrocities. For diffusion over smaller groups (for example, neighborhoods) in which network characteristics are critical for understanding the spread or prevention of mass atrocities, much more is to be learned about how

small (even trivial) “trembles” on the network can set off a contagion of mass atrocity acceptance or to stop the contagion before it takes off. For models of tentacles of atrocity perpetration, other “market structures” can be modeled including Stackelberg-like behavior in which the atrocity perpetrators’ harming actions are given and a third party optimally allocates efforts on the network to maximize the saving of lives. Similarly, the set of reaction functions for both the atrocity actors and the third party intervenors can be modeled to find the Cournot-Nash equilibrium. The CSF on networks approach can be further refined by assuming that the competing networks allocate their efforts to particular locales of contestation in which each locale has its own CSF and the leaders of each network are strategic in their decision-making. Furthermore, the contesting networks approach to modeling atrocity onset, spread, and prevention is amenable to analysis of resistance actors (saboteurs) at key places in the network.

References

- Acemoglu, Daron, Azarakhsh Malekian, and Asu Ozdaglar. 2016. “Network Security and Contagion,” *Journal of Economic Theory* 166(November): 536–85.
- Adalian, Rouben Paul. 2009. “The Armenian Genocide.” In Samuel Totten and William S. Parsons, eds. *Century of Genocide: Critical Essays and Eyewitness Accounts*. Third Edition, 55–92. New York: Routledge.
- Anderton, Charles H., and Jurgen Brauer, eds. 2018. *Economic Aspects of Genocides, Other Mass Atrocities, and Their Prevention*. New York: Oxford University Press.
- Anderton, Charles H., and Jurgen Brauer. 2018. “Mass Atrocities and Their Prevention.” Working Paper (not posted).

- Baccara, MariaGiovanna, and Heski Bar-Isaac. 2008. "How to Organize Crime." *Review of Economic Studies* 75: 1039–67.
- Ballester, Coralio, Antoni Calvó-Armengol, and Yves Zenou. 2006. "Who's Who in Networks. Wanted: The Key Player." *Econometrica* 74(5): 1403–17.
- Bass, Frank. 1969. "A New Product Growth for Model Consumer Durables." *Management Science* 15 (5): 215–27.
- Bonacich, Phillip. 1987. "Power and Centrality: A Family of Measures." *American Journal of Sociology* 92(5): 1170–82.
- Bramoullé, Yann, Andrea Galeotti, and Brian W. Rogers, eds. 2016. *The Oxford Handbook of the Economics of Networks*. New York: Oxford University Press.
- Browning, Christopher R. 1992. *The Path to Genocide: Essays on Launching the Final Solution*. New York: Cambridge University Press.
- Browning, Christopher R. 2004. *The Origins of the Final Solution: The Evolution of Nazi Jewish Policy, September 1939–March 1942*. Lincoln, NE: University of Nebraska Press.
- Campbell, Bradley. 2010. "Contradictory Behavior during Genocides." *Sociological Forum* 25(2): 296–314.
- Cunningham, William H. 1985. "Optimal Attack and Reinforcement of a Network." *Journal of the Association for Computing Machinery* 32(3): 549–61.
- Des Forges, Alison. 1999. *Leave None to Tell the Story: Genocide in Rwanda*. Washington, D.C.: Human Rights Watch.
- Donà, Giorgia. 2018. "'Situated Bystandership' During and After the Rwandan Genocide." *Journal of Genocide Research* 20(1): 1–19.

- Dziubiński, Marcin, Sanjeev Goyal, and Adrien Vigier. 2016. "Conflict and Networks." In *The Oxford Handbook of the Economics of Networks*, edited by Yann Bramoullé, Andrea Galeotti, and Brian Rogers, 215–43. New York: Oxford University Press.
- Easley, David, and Jon Kleinberg. 2010. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. New York: Cambridge University Press.
- Enders, Walter, and Todd Sandler. 2012. *The Political Economy of Terrorism, 2nd edition*. New York: Cambridge University Press.
- Franke, Jörg, and Tahir Öztürk. 2009. "Conflict Networks." Working paper 116. Dortmund, Germany: Technische Universität Dortmund.
- Fujii, Lee A. 2009. *Killing Neighbors: Webs of Violence in Rwanda*. Ithaca, NY: Cornell University Press.
- Gintis, Herbert. 2009. *Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences*. Princeton, NJ: Princeton University Press.
- Gobodo-Madikizela, Pumla. 2002. "Remorse, Forgiveness, and Rehumanization: Stories from South Africa." *Journal of Humanistic Psychology* 42(1): 7–32.
- Goldhagen, Daniel J. 1996. *Hitler's Willing Executioners: Ordinary Germans and the Holocaust*. New York: Knopf.
- Goyal, Sanjeev, and Adrien Vigier. 2014. "Attack, Defence, and Contagion in Networks." *Review of Economic Studies* 81: 1518–42.
- Granovetter, Mark. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology* 83: 489-515.
- Jackson, Matthew O. 2008. *Social and Economic Networks*. Princeton, NJ: Princeton University Press.

- Kleinschmidt, Horst, ed. 1972. *White Liberation*. Johannesburg: Spro-cas.
- König, Michael D., Dominic Rohner, Matthias Thoenig, and Fabrizio Zilibotti. 2017. "Networks in Conflict: Theory and Evidence from the Great War of Africa." *Econometrica* 85(4): 1093–132.
- Lamberson, P.J. 2016. "Diffusion in Networks." In Yann Bramoullé, Andrea Galeotti, and Brian W. Rogers, eds. *The Oxford Handbook of the Economics of Networks*, 479–503. New York: Oxford University Press.
- Lemarchand, René. 2009. "The 1994 Rwandan Genocide." in Samuel Totten and William S. Parsons, eds. *Century of Genocide: Critical Essays and Eyewitness Accounts*. Third Edition, 483–504. New York: Routledge.
- McDoom, Omar Shahabudin. 2013. "Who Killed in Rwanda's Genocide? Micro-space, Social Influence and Individual Participation in Intergroup Violence," *Journal of Peace Research* 50(4): 453–67.
- McDoom, Omar Shahabudin. 2014. "Predicting Violence within Genocide: A Model of Elite Competition and Ethnic Segregation from Rwanda." *Political Geography* 42 (September): 34–45.
- Michaeli, Moti, and Daniel Spiro. 2017. "From Peer Pressure to Biased Norms." *American Economic Journal: Microeconomics* 9(1): 152–216.
- Maoz, Zeev. 2011. *Networks of Nations: The Evolution, Structure, and Impact of Internationalized Networks, 1816-2011*. New York: Cambridge University Press.
- Newman, Mark E.J. 2010. *Networks: An Introduction*. New York: Oxford University Press.
- Petrova, Maria, and David Yanagizawa-Drott. 2016. "Media Persuasion, Ethnic Hatred, and Mass Violence: A Brief Overview of Recent Research Advances." In *Economic Aspects of*

- Genocides, Other Mass Atrocities, and Their Prevention*, edited by Charles H. Anderton and Jurgen Brauer, 274–86. New York: Oxford University Press.
- Ray, Subhasish. 2017. “Intra-Group Interactions and Inter-Group Violence: Sikh Mobilization during the Partition of India in a Comparative Perspective.” *Journal of Genocide Research* 19(3): 382–403.
- Roth, Paul A. 2010. “Social Psychology and Genocide.” In Donald Bloxham and A. Dirk Moses, eds., *The Oxford Handbook of Genocide Studies*, 198–216. New York: Oxford University Press.
- Scaife, Laura. 2017. *Social Networks as the New Frontier of Terrorism: #Terror*. New York: Routledge.
- Sen, Amartya K. 2006. *Identity and Violence: The Illusion of Destiny*. London: Allen Lane.
- Schelling, Thomas C. 1971. “Dynamic Models of Segregation.” *Journal of Mathematical Sociology* 1: 143–86.
- Schelling, Thomas C. 1978. *Micromotives and Macrobehavior*. New York: W.W. Norton.
- Staub, Ervin. 2005. “Constructive Rather Than Harmful Forgiveness, Reconciliation, and Ways to Promote them after Genocide and Mass Killing.” In Everett L. Worthington, Jr., ed., *Handbook of Forgiveness*, 443–59. New York: Routledge.
- Straus, Scott. 2006. *The Order of Genocide*. Ithaca, NY: Cornell University Press.
- Terrorism Network Project. 2018. “Terrorism Network Project.” Available: https://www.du.edu/korbel/sie/research/chenow_terrorism_networks.html. Accessed: September 10, 2018.
- Valentino, Benjamin. 2004. *Final Solutions: Mass Killing and Genocide in the 20th Century*. Ithaca, NY: Cornell University Press.

- Waller, James. 2007. *Becoming Evil: How Ordinary People Commit Genocide and Mass Killing*. Second Edition. New York: Oxford University Press.
- Wilson, Paul. 2010. "Do Normal People Commit Genocide? Observations from the Cambodian Trial of 'Duch'." *Psychiatry, Psychology and Law* 17(4): 495–502.
- Yanagizawa-Drott, David. 2014. "Propaganda and Conflict: Evidence from the Rwandan Genocide." *Quarterly Journal of Economics* 129(4): 1947-94.
- Zech, Steven T., and Michael Gabbay. 2016. "Social Network Analysis in the Study of Terrorism and Insurgency: From Organization to Politics." *International Studies Review* 18(2): 214–43.

Figure 1




Figure 2




Figure 3

Intrinsic reservation value, $r(x)$,
 Cost of atrocity perpetration, c




Figure 4

Reservation valuations
 and cost




Figure 5




Figure 6




Figure 7




Figure 8

		Player 2	
		A	P
Player 1	A	a, a	$0, 0$
	P	$0, 0$	b, b

Figure 9








Figure 10

Figure 11




Figure 12




Figure 13







Figure 14




Figure 15

