Social Norms and Social Change*

Ethan Bueno de Mesquita† Mehdi Shadmehr‡

August 21, 2020

Abstract

Applying the literature on beauty contest models, we study how social norms affect social change in a dynamic setting. We apply our model to two contemporary crises: police misconduct and social distancing following COVID-19. Because people have a desire to socially conform and past experience helps coordinate behavior, aggregate behavior has inherent inertia. In the COVID-19 application, this implies that aggregate social distancing lags behind the social optimum after the disease outbreak. In the policing application, it implies high levels of misconduct are expected to persist even if the ubiquity of cell phone video, shifting public opinion, or reforms to recruitment or training succeed in changing officers' attitudes. Public statements by leaders facilitate coordination on new norms, reducing inertia and facilitating social change. In the policing application, public messages are preferable to private messages. In the COVID-19 application, public communication is preferable when optimal social distancing is highly correlated over time and when individuals are poorly-informed about changes in optimal social distancing.

*We have received helpful comments from Scott Ashworth, Dan Bernhardt, Wioletta Dziuda, Amanda Friedenberg, and Arda Gitmez. We thank the Becker Friedman Institute for financial support.

†Harris School of Public Policy, University of Chicago. Email:bdm@uchicago.edu

‡Department of Economics, University of Calgary. E-mail: mehdi@uchicago.edu

In many settings, people's preferences depend on both internal motivations and on a desire to conform to the behavior of others. An important literature shows that in such settings, when people are also uncertain about one another's internal motivations, public information shapes social norms. In particular, individual decision makers overweight public information because it plays a coordinating role, helping people conform to aggregate behavior (Morris and Shin 2002).

We study a dynamic version of the canonical model in order to ask how such norm-based behavior affects the possibility of social change. In our model, people have both internal and conformist motivations. They are uncertain about others' internal motivations. And internal motivations are changing over time, creating the possibility of desirable social change.

We relate this model to two contemporary crises: COVID-19 and police misconduct.

In the case of COVID-19, we relate our model to decisions about social distancing. Individuals are motivated to socially distance in a way that benefits society. (Jordan et al 2020 provide evidence that people have pro-social preferences concerning social distancing.) However, the socially optimal level of distancing shifts over time and is uncertain, creating uncertainty about how others will behave. Moreover, individuals also want their behavior to conform to others' behavior. People feel awkward wearing a mask, risking offending a neighbor by waiting for the next elevator, declining an outstretched hand, or forbidding playdates, if others are not doing likewise.

In the case of policing, we relate our model to officers' decisions about violent confrontations, discrimination, or support for other officers who have engaged in misconduct. Officers have their own individual views on the use of physical force, discrimination, or support for others who have engaged in misconduct. Such attitudes may be changing over time—for instance, due to the ubiquity of cell phone video, changing views of voters or politicians, or improved selection and training of officers—and their is uncertainty about the attitudes of other officers. Moreover, individual police officers also want to conform to the behavior of their colleagues, so as to be perceived as team players.

Our first result shows that aggregate behavior exhibits inertia (Proposition 2). Individual's actions overweight their common knowledge about past behavior. This is because adhering to past norms helps to coordinate behavior, facilitating conformity with the behavior of others. As a result of such social norms, changes in social behavior lag behind changes in internal motivations.

In the case of COVID-19, this result implies that, following a large positive shock to the socially optimal level of social distancing (e.g., due to the outbreak of a contagious virus), aggregate behavior will be far below the social optimum, even if individuals' information accurately reflects this shock on average.

Importantly, in the case of policing, the model ought not be interpreted in the same normative terms as in the application to COVID-19. In particular, the normative goal should surely be to minimize violent confrontation, discrimination, and support for misconduct. In this context, the result shows that even if officers become individually less inclined to use force, discriminate, or support misconduct, aggregate policing behavior will lag behind.

How can the government improve this situation? Our Proposition 3 suggests that public statements from prominent leaders can reduce inertia by generating new common knowledge that enables people to coordinate their actions around new, more appropriate norms. Should such statements be made publicly or privately? In the case of COVID-19, our Proposition 5 suggests that public communication is better than private communication when optimal social distancing behavior is highly correlated over-time and when individuals are poorly-informed, so that the overweighting of prior behavioral norms is more severe. This result reflects that the information communicated publicly creates new common knowledge. As such, people concerned with conforming to social norms overweight the new public information. This overweighting can be beneficial because it helps offset the overweighting of past behavioral norms.

We utilize a standard "beauty contest" model in which individuals aim to take actions that match an uncertain state of the world while remaining close to the average action in the population. Morris and Shin's (2002) seminal paper showed that individuals overreact to public information in such settings. Subsequent literature explored the welfare consequences of this insight and its implications for optimal communication (Morris and Shin 2007; Angeletos and Pavan 2007), as well as its extensions to dynamic settings (Morris and Shin 2006; Angeletos and La'O 2010; Huo and Pedroni 2020). This framework has been applied to study a wide range of topics, including leadership in party conferences (Dewan and Myatt 2008), organizations (Bolton et al. 2013; Landa and Tyson 2017), and judiciaries (Shadmehr et al. 2019). Our paper provides an application to social change, with an emphasis on contemporary social and policy crises.

Our paper is also related to the game theoretic literature on social norms. Much of this literature interprets social norms as a set of equilibrium expectations and behaviors, so that different social norms correspond to different equilibria of a game with multiple equilibria (Myerson 1991, p. 113-114; Postlewaite 2011; Bidner and Franscois 2013; Acemoglu and Jackson 2015; Young 2015). Chwe (1998, 2013) argues that many social interactions and rituals exist to create common knowledge that coordinates behavior on some social norm. Such rituals are common in organizations like the police, where long standing institutional practices, a shared language and mythology, and a variety of rites of passage create powerful cultural expectations. Indeed, a long sociological tradition explores the role of social norms in shaping policing culture (Wilson 1978, Moskos 2008, Patterson 2014, Braga et al. 2019). Our paper shows that the common knowledge created by such practices make social change difficult.

We adopt Acemoglu and Jackson's (2017) definition of social norms as "the distribution of anticipated payoff-relevant behavior" (p. 246). Our model has a unique equilibrium, and social norms in our setting refer to the average behavior of the population. When agents are heterogeneously informed, the anticipated average behavior depends on the individual's information. We model people's desire to conform to social norms as a desire to do what is expected of them, as reflected in other people's behavior.

Acemoglu and Jackson (2015) study norms of cooperation in a dynamic setting where an agent in each period plays a complete information coordination game with agents of the immediate past and future periods. Their model features behavioral types who never cooperate and past actions are observed with noise, so some rational agents do not cooperate even in the best equilibrium. When the cooperative action of a generation becomes public for future generations, their expectations about cooperative behavior improves, facilitating cooperation—see also Acemoglu and Wolitsky (2014) and Acemoglu and Jackson (2017). In contrast to these papers, agents in our setting are uncertain about one anothers' private motivations and about one anothers' information in a changing world. While Acemoglu and Jackson (2015, 2017) focus on patterns of cooperation and compliance, our focus is on inertia and over-reactions to public information in a changing world and on the resulting policy implications.

1 Model

We apply the canonical framework and results discussed in Angeletos and Lian's (2016) comprehensive review of the literature.[1] There is a continuum of individuals indexed by $i \in [0, 1]$, interacting over time, indexed by $t = 0, 1, \ldots$. In each period t, each individual must take an action $a_{it} \in \mathbb{R}$. In the COVID-19 application, a higher action corresponds to more social distancing by the individual. In the policing application, a higher action corresponds to more violent confrontations, discrimination, or support for other officers who have engaged in misconduct.

Absent concerns for conformity, an individual's preferred action is θ_t. We refer to θ_t as the *target action* in period t. In the COVID-19 application, we interpret this target action as corresponding to the socially optimal level of social distancing. In the policing application, we interpret the target action as corresponding to police officers' views about the level of violent confrontation, discrimination, or support for misconduct they should engage in.

But, in each period, each individual cares about this target action *and* about conforming to the average action that others take in that period, $A_t = \int a_{it}\, di$. This generates a complementarity: if an individual believes that others do little social distancing or a lot of misconduct, this raises that individual's incentive also to do less social distancing or more misconduct. This captures, among other things, social pressure and the cost of deviating from the norms of behavior in the society.

[1] This literature builds on the seminal work of Morris and Shin (2002).

A individual's payoff in period t is:

$$-(1-\alpha)(a_{it}-\theta_t)^2-\alpha(a_{it}-A_t)^2, \tag{1}$$

where $\alpha \in (0,1)$ is the individual's relative weight on conformity.

The target action, θ_t, follows a random walk: $\theta_t = \theta_{t-1} + u_t$, where $u_t \sim iidN(0, \sigma_u)$. Individual's do not observe θ_t, but each individual observes a signal: $x_{it} = \theta_t + \epsilon_{it}$, where $\epsilon_{it} \sim iidN(0, \sigma_\epsilon)$. Throughout, we assume that the noise and fundamentals are independent from each other in the standard manner. So people are uncertain about one another's beliefs about the target action.

Individual i observes x_{it} in period t, and θ_{t-1} becomes public in period t.[2] Individuals discount future payoffs by δ, and each individual maximizes the expected sum of discounted period payoffs.

2 Analysis

Because there is a continuum of individuals, an individual's action does not affect the aggregate outcome, either in the current or in future periods. Thus, the only link between periods is information. From equation (1), individual i chooses the following action:

$$a_{it} = (1-\alpha)\mathbb{E}_{it}[\theta_t] + \alpha\ \mathbb{E}_{it}[A_t], \tag{2}$$

where $\mathbb{E}_{it}[\cdot]$ is the expectation of i in period t given his information.

Define $\bar{\mathbb{E}}^h$ recursively as follows. $\bar{\mathbb{E}}^0[X] = X$, $\bar{\mathbb{E}}^1[X] = \bar{\mathbb{E}}[\bar{\mathbb{E}}^0[X]] = \int \mathbb{E}_i[X]di$, $\bar{\mathbb{E}}^h[X] = \bar{\mathbb{E}}[\bar{\mathbb{E}}^{h-1}[X]] = \int \mathbb{E}_i[\bar{\mathbb{E}}^{h-1}[X]]di$. That is, $\bar{\mathbb{E}}^1[X]$ is the average expectation of the random variable X in the population; $\bar{\mathbb{E}}^2[X]$ is the average expectation in the population about the average expectation in the population, and so on. Proposition 1 shows that the aggregate action in the population depends on all such higher order expectations in the population about the target action. The proof comes from iterating on equation (2). (All proofs are in the appendix.)

Proposition 1 *The aggregate action in each period depends on all average higher order beliefs in the population about the target action, with lower weights on higher orders:*

$$A_t = \sum_{h=1}^{\infty}(1-\alpha)\ \alpha^{h-1}\ \bar{\mathbb{E}}_t^h[\theta_t] \tag{3}$$

Observe that

$$\mathbb{E}_{it}[\theta_t] = \mathbb{E}_{it}[u_t] + \theta_{t-1} \ \Rightarrow\ \bar{\mathbb{E}}_t^h[\theta_t] = \bar{\mathbb{E}}_t^h[u_t] + \theta_{t-1}. \tag{4}$$

[2] Even if θ_{t-1} is not observed in the current period, individuals will infer it in equilibrium if they observe the last period's aggregate behavior A_{t-1}.

Now, using properties of Normal distribution and Proposition 1 yields:

Proposition 2 *Conformity generates inertia. In particular,*

- $A_t = \theta_{t-1} + \phi u_t$, *where* $0 < \phi < \beta < 1$, $\phi = \frac{(1-\alpha)\beta}{1-\alpha\beta}$, *and* $\beta = \frac{\sigma_u^2}{\sigma_u^2+\sigma_\epsilon^2}$.
- ϕ *is decreasing in* α, *with* $\lim_{\alpha\to 0}\phi(\alpha) = \beta$ *and* $\lim_{\alpha\to 1}\phi(\alpha) = 0$.

Because individuals care about coordinating their actions, they put extra weight on their common knowledge of the past, which facilitates coordination—reminiscent of the logic of focal points. Individuals have common knowledge that, on average, today's target action is yesterday's target action (i.e., $\theta_t \sim N(\theta_{t-1}, \sigma_u)$) and hence over-weight this fact. As a result, today's aggregate action is biased in the direction of yesterday's target action.

In our COVID-19 application, Proposition 2 implies that, following a large positive shock to the socially optimal amount of social distancing (like the outbreak of COVID-19), aggregate social distancing will be far lower than is socially optimal.[3] And in our policing application, even if features of the environment—like the ubiquity of cell phones, changing views among voters and politicians, or reform of the selection or training process for officers—make individual police officers desire to engage in misconduct, such behavior will persist because past misconduct has created expectations of future misconduct.

Now, suppose in each period t, in addition to the their private signals x_{it}, individuals also receive a public signal $p_t = \theta_t + \eta_t$, with $\eta_t \sim iidN(0, \sigma_\eta)$. Such a public signal might be the result of information conveyed by leaders. Proposition 1 and equation (4) still hold because they do not depend on the details of available information. However, the presence of public signals changes the degree of inertia in aggregate behavior.

Proposition 3

1. *Public signals reduce inertia. Averaging over the public signal noise, the expected aggregate action is:* $E[A_t|\theta_{t-1}, u_t] = \theta_{t-1} + \phi_p u_t$, *where* $\phi < \phi_p < 1$. *Moreover,* $\phi_p - \phi$ *is increasing in* α.
2. *The amount of inertia is decreasing in the clarity of the public signal. That is,* ϕ_p *is monotone decreasing in* σ_η, $\lim_{\sigma_\eta\to\infty}\phi_p(\sigma_\eta) = \phi$, *and* $\lim_{\sigma_\eta\to 0}\phi_p(\sigma_\eta) = 1$.

3 Normative Implications

In the policing application, the natural normative goal is to minimize violence encounters, discrimination, and support of misconduct. As such, an immediate implication of Proposition 3 is that, in periods in which θ decreases—i.e., police themselves are coming to view

[3]The over-weighting of public information is a key insight of Morris and Shin (2002) who showed that public announcements can be damaging to welfare, particularly, in financial settings.

such behavior as less desirable—public communication about expected behavior will improve matters by reducing inertia. Moreover, fixing the amount of information conveyed, inertia is reduced more when information is conveyed publicly versus privately.

Matters are more subtle in the COVID-19 application. First, in this setting, complete social distancing is not the right normative benchmark. Rather, we will think of the normative goal as each individual choosing the target action, $a_{it} = \theta_t$, so that the aggregate level of social distancing is socially optimal, $A_t = \int a_{it} di = \theta_t$. In light of this, when discussing social distancing, we say that any reduction in the expected quadratic distance between individual actions and the right action, $E[\int (a_{it} - \theta_t)^2 di]$, is a *normative improvement.*

Proposition 3 shows that, following a shock, an informed leader can send a public signal about the target action that helps set public expectations, thereby reducing inertia driven by the desire to conform. The clearer that message (i.e., the lower σ_η), the more this will reduce inertia. The next result shows when this is a normative improvement.

Proposition 4 *Improving the clarity of the public signal causes a normative improvement in social distancing (i.e., $E[\int (a_{it} - \theta_t)^2 di]$ is increasing in σ_η) if: (i) $\alpha \leq 1/2$ or (ii) σ_η is sufficiently small.*

In the social distancing case, public messages are a normative improvement if people don't put too much weight on conformity ($\alpha \leq 1/2$) or the public signal is sufficiently informative (σ_η small). Why these conditions? Because individuals value conformity, they put excessive weight on all public signals relative to a Bayesian individual who only cares about choosing an action that reflects the best estimate of θ_t (this was the same logic that drove inertia in the first place). Because this distortion is smaller when α is smaller, new public information about the optimal social distancing is always beneficial when citizens put relatively less weight on conformity ($\alpha \leq 1/2$). In the other extreme, when individuals almost only care about conformity ($\alpha \approx 1$), they put almost no weight on their private signals. Now, although individuals over-weight new public information (p_t), this over-reaction to the new public information helps counter-act their over-reaction to past experience (that $\theta_t \sim (\theta_{t-1}, \sigma_u^2)$), and the overall effect is again beneficial. In between, when $\alpha \in (1/2, 1)$, these effects compete and the overall effect of raising the precision of new public information may be negative unless it is sufficiently informative (σ_η small) to offset the over-reaction.[4]

For social distancing in the presence of a dangerous infectious disease, we believe the relevant parameter space is $\alpha \leq 1/2$. It is unlikely that people care so much about conformity that over-reaction to new public information trumps its value. Hence, for cases like COVID-19, Proposition 3 suggests that clear and consistent public messages from a leader are likely to be socially beneficial.

Given the overreaction by individuals to public messages described above, one may wonder whether there is a better way to deliver information. Would it be better for individuals

[4]Equation (16) in the proof of Proposition 3 shows the necessary and sufficient conditions for when reducing σ_η is a normative improvement.

to receive the same level of information, but privately rather than publicly? For instance, perhaps employers or local governments could provide private information to individuals, rather than them all observing the same public information in a presidential speech or press conference.

To consider this possibility, contrast the public signal case with a setting where, instead of receiving private and public signals $x_{it} \sim N(\theta_t, \sigma_\epsilon^2)$ and $p_t \sim N(\theta_t, \sigma_\eta^2)$, citizens receive a single private signal x'_{it} with the same amount of information about the right action θ_t as the public and private signals combined. In particular, let $x'_{it} = \theta_t + \epsilon'_{it}$, with $\epsilon'_{it} \sim N(0, \sigma_{\epsilon'}^2 = \frac{\sigma_\epsilon^2 \sigma_\eta^2}{\sigma_\epsilon^2 + \sigma_\eta^2})$.

Proposition 5 *The setting with the combination of private and public signals* (x_{it}, p_t) *is a normative improvement over the setting with more precise private signals* x'_{it} *when* σ_u *is sufficiently small or* σ_ϵ *is sufficiently large.*

When individuals believe the past is highly informative about the present (σ_u small) or that they are privately poorly-informed (σ_ϵ large), individuals put too much weight on their past experience. In such circumstances, it is better for the government to communicate publicly rather than privately. Individuals over-react to the government's public messages. But that will help to counter-act their over-reaction to their past experience. By contrast, when individuals believe the past is relatively uninformative (σ_u large) or that they are privately well-informed (σ_ϵ small), the government should communicate privately.

Standard accounts frame the problem of social distancing as a public goods problem with the familiar externalities. For example, individuals do not internalize that social distancing has positive health spillovers on others, so there will be under-provision of social distancing. In this setting, the fundamental problem is not informational: even fully-informed citizens under-provide social distancing absent some more heavy-handed policy that directly changes citizen incentives—e.g., forced downtown closure or fines for public gatherings. Moreover, enforcing behavioral changes for actions that are largely taken out of the public eye, such as hand washing, handshakes, or private gatherings, is virtually impossible. This aspect of the social distancing challenge has been the focus of public and academic discussions. For example, Allcott et al. (2020) study the interaction between risk perception and such externalities in the United States. Dube and Baicker (2020) discuss the importance of individuals sacrificing their interests for the greater good and, drawing on Christensen et al.'s (2020) study of the Ebola crisis in Sierra Leon, emphasize the importance of trust in local leaders and institutions.

In contrast, our analysis highlights the role of social norms and strategic uncertainty as an under-appreciated source of counter-productive inertia in aggregate social distancing. The policy implications are also sharply different. While information alone cannot resolve the standard externalities problem, clear and consistent *public* information can dramatically improve social distancing by reducing strategic uncertainty and enabling citizens to coordinate on new norms. Moreover, in contrast to accounts that emphasize the role of the local community in providing trusted information, we highlight the advantage of information provided by national over local leaders, especially in countries like United States where there is trust in the expertise of governmental health organizations. National coverage generates more common knowledge, enabling citizens to better coordinate on new optimal norms of social distancing.

4 Conclusion

Social norms can disrupt valuable social change. In particular, aggregate behavior overreacts to past experience when people are motivated to socially conform and are uncertain about one another's beliefs. This is because the common knowledge created by past experience creates social norms that facilitate coordination. The result is undesirable inertia—behavior tends to fall back toward old norms. In the case of COVID-19, this suggests that there is too little social distancing following disease outbreaks or flare ups. In the case of police misconduct, it suggests that, even if police attitudes shift in a positive direction, police behavior will lag behind.

Communication by leaders can help to mitigate this damaging inertia. If messages are public, the common knowledge they create engenders an over-reaction analogous to that created by past experience. This over-reaction can help off-set the inertia resulting from the over-weighting of past norms. In the case of police misconduct, such communication is beneficial as long as personal motivations are in fact trending toward a preference for less violence, discrimination, or support for misconduct. In the case of the outbreak of an infectious disease, public communication is preferable to private communication if individuals' information is noisy (σ_ϵ is large) and the optimal level of social distancing is sticky (σ_u is low, the disease is a very unusual shock). In the case of COVID-19, both of these conditions are likely to hold. As such, clear and consistent public statements by a national leader are expected to be more effective than similarly informative, but more private, communications (e.g., by local governments or employers).

5 References

Acemoglu, Daron, and Matthew O. Jackson. 2015. "History, Expectations, and Leadership in the Evolution of Social Norms." *Review of Economic Studies* 82: 1-34.

Acemoglu, Daron, and Matthew O. Jackson. 2017. "Social Norms and the Enforcement of Laws." *Journal of the European Economic Association* 12: 245-95.

Acemoglu, Daron, and Alexander Wolitzky. 2014. "Cycles of Conflict: An Economic Model." *American Economic Review* 104: 1350-67.

Allcott, Hunt, Levi Boxell, Jacob Conway, Matthew Gentzkow, Michael Thaler, and David

Yang. 2020. "Polarization and Public Health: Partisan Differences in Social Distancing during the Coronavirus Pandemic." Mimeo.

Angeletos, George-Marios, and Alessandro Pavan. 2007. "Efficient Use of Information and Social Value of Information." *Econometrica* 75: 1103–42.

Angeletos, George-Marios, and Jennifer La'O. 2010. "Noisy Business Cycles." in *NBER Macroeconomics Annual 2009, Volume 24*, pp. 319—378, University of Chicago Press.

Angeletos, George-Marios, and Chen Lian. 2016. "Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination." *Handbook of Macroeconomics, Vol. 2*: 1065–1240.

Bidner, Chris, and Patrick Francois. 2013. "The Emergence of Political Accountability." *Quarterly Journal of Economics* 128: 1397–1448.

Bolton, Patrick, Markus K. Brunnermeier, and Laura Veldkamp. 2013. "Leadership, Coordination and Corporate Culture." *Review of Economic Studies* 80: 512–37.

Braga, Anthony A., Rod K. Brunson, and Kevin M. Drakulich. 2019. "Race, Place, and Effective Policing." *Annual Review of Sociology* 45:535–555.

Christensen, Darin, Oeindrila Dube, Johannes Haushofer, Bilal Siddiqi, and Maarten Voors. 2020. "'Building Resilient Health Systems: Experimental Evidence from Sierra Leone and the 2014 Ebola Outbreak." Mimeo.

Chwe, Michael Suk-Young. 1998. "Culture, Circles, and Commercials: Publicity, Common Knowledge, and Social Coordination." *Rationality and Society* 10(1):47–75.

Chwe, Michael Suk-Young. 2013. *Rational Ritual: Culture, Coordination, and Common Knowledge.* Princeton University Press.

Dewan, Torun, and David P. Myatt. 2008. "The Qualities of Leadership: Direction, Communication, and Obfuscation." *American Political Science Review* 102: 351-68.

Dube, Oeindrila, and Katherine Baicker. 2020. "How You Can Protect Your Community, Not Just Your Own Health." *New York Times, March 26, 2020.*

Huo, Zhen, and Marcelo Pedroni. 2020. "A Single-Judge Solution to Beauty Contests." *American Economic Review* 110: 526-68.

Jordan, Jillian J., Erez Yoeli, and David G. Rand. 2020. "Don't get it or don't spread it? Comparing self-interested versus prosocially framed COVID-19 prevention messaging" PsyArXiv, 3 April. `https://doi.org/10.31234/osf.io/yuq7x`

Landa, Dimitri, and Scott Tyson. 2017. "Coercive Leadership." *American Journal of Political Science* 61: 559-74.

Morris, Stephen, and Hyun Song Shin. 2002. "The Social Value of Public Information." *American Economic Review* 92: 1521–34.

Morris, Stephen, and Hyun Song Shin. 2006. "Inertia of Forward-Looking Expectations." *AEA Papers & Proceedings* 96: 152–7.

Morris, Stephen, and Hyun Song Shin. 2007. "Optimal Communication." *Journal of the European Economic Association* 5: 594–602.

Moskos, Peter. 2008. *Cop in the Hood: My Year Policing Baltimore's Eastern District.* Princeton: Princeton University Press.

Myerson, Roger. 1991. *Game Theory: Analysis of Conflict.* Boston, MA: Harvard University Press.

Patterson, Orlando. 2014. "Making Sense of Culture." *Annual Review of Sociology* 40: 1–30.

Postlewaite, Andrew. 2011. "Social Norms and Preferences." *Handbook of Social Economics, Volume 1A,* edited by J. Benhabib, A. Bisin, and M. Jackson. Amsterdam: North-Holland.

Shadmehr, Mehdi, Charles Cameron, and Sepehr Shahshahani. 2019. "Coordination and Innovation in Judiciaries: Correct Law vs. Consistent Law." Mimeo.

Wilson, James Q. 1978. *Varieties of Police Behavior: The Management of Law and Order in Eight Communities.* Cambridge: Harvard University Press.

Young, Peyton. 2015. "The Evolution of Social Norms." *Annual Review of Economics* 7: 359–87.

6 Appendix: Proofs

Proof of Proposition 1: From equation (2),

$$A_t = \int a_{it} di = \int ((1-\alpha)\mathbb{E}_{it}[\theta_t] + \alpha\ \mathbb{E}_{it}[A_t])di = (1-\alpha)\bar{\mathbb{E}}_t[\theta_t] + \alpha\ \bar{\mathbb{E}}_t[A_t].$$

Iterating yields:

$$A_t = (1-\alpha)\bar{\mathbb{E}}_t[\theta_t] + (1-\alpha)\alpha\ \bar{\mathbb{E}}_t^2[\theta_t] + \alpha^2\bar{\mathbb{E}}_t^2[A_t].$$

Repeated iteration yields the result. □

Proof of Proposition 2: We calculate $\bar{\mathbb{E}}_t^h[\theta_t]$, and use Proposition 1. Note that $x_{it} = \theta_t + \epsilon_{it} = \theta_{t-1} + u_t + \epsilon_{it}$. Thus, letting $\beta = \frac{\sigma_u^2}{\sigma_u^2+\sigma_\epsilon^2}$

$$\mathbb{E}_{it}[u_t] = \beta(x_{it} - \theta_{t-1}) = \beta(u_t + \epsilon_{it}) \ \Rightarrow\ \bar{\mathbb{E}}_t[u_t] = \beta u_t.$$

Iterating yields:

$$\bar{\mathbb{E}}_t^h[u_t] = \beta^h\ u_t. \tag{5}$$

Substituting from equation (5) into equation (4) yields:

$$\bar{\mathbb{E}}_t^h[\theta_t] = \beta^h \ u_t + \theta_{t-1}. \tag{6}$$

Now, substituting from equation (6) into equation (3) in Proposition 1 yields:

$$A_t = \sum_{h=1}^{\infty}(1-\alpha)\ \alpha^{h-1}\ (\beta^h\ u_t + \theta_{t-1}) = \theta_{t-1} + \frac{\beta(1-\alpha)}{1-\alpha\beta}\ u_t. \tag{7}$$

In equation (7), let $\phi = \frac{(1-\alpha)\beta}{1-\alpha\beta}$, and observe that $\lim_{\sigma_\epsilon\to 0}\beta = \lim_{\sigma_u\to\infty}\beta = 1$. $\square$

Proof of Proposition 3: With the public signal p_t, $\mathbb{E}_{it}[u_t] = E[u_t|x_{it}, p_t]$. Thus,

$$\begin{aligned}\mathbb{E}_{it}[u_t] &= \frac{\sigma_u^2\sigma_\eta^2(x_{it}-\theta_{t-1}) + \sigma_u^2\sigma_\epsilon^2(p_t-\theta_{t-1})}{\sigma_u^2\sigma_\eta^2 + \sigma_u^2\sigma_\epsilon^2 + \sigma_\epsilon^2\sigma_\eta^2} \\ &= \frac{\sigma_u^2\sigma_\eta^2(u_t+\epsilon_{it}) + \sigma_u^2\sigma_\epsilon^2(p_t-\theta_{t-1})}{\sigma_u^2\sigma_\eta^2 + \sigma_u^2\sigma_\epsilon^2 + \sigma_\epsilon^2\sigma_\eta^2}.\end{aligned} \tag{8}$$

Thus,

$$\bar{\mathbb{E}}_t[u_t] = \frac{\sigma_u^2\sigma_\eta^2 u_t + \sigma_u^2\sigma_\epsilon^2(p_t-\theta_{t-1})}{\sigma_u^2\sigma_\eta^2 + \sigma_u^2\sigma_\epsilon^2 + \sigma_\epsilon^2\sigma_\eta^2} = A_u\ u_t + A_p\ (p_t - \theta_{t-1}), \tag{9}$$

where

$$A_u = \frac{\sigma_u^2\sigma_\eta^2}{\sigma_u^2\sigma_\eta^2 + \sigma_u^2\sigma_\epsilon^2 + \sigma_\epsilon^2\sigma_\eta^2} \quad \text{and} \quad A_p = \frac{\sigma_u^2\sigma_\epsilon^2}{\sigma_u^2\sigma_\eta^2 + \sigma_u^2\sigma_\epsilon^2 + \sigma_\epsilon^2\sigma_\eta^2}, \tag{10}$$

with

$$\lim_{\sigma_\eta\to\infty} A_p = 0,\ \lim_{\sigma_\eta\to\infty} A_u = \beta,\ \lim_{\sigma_\eta\to 0} A_p = 1,\ \text{and}\ \lim_{\sigma_\eta\to 0} A_u = 0. \tag{11}$$

Iterating on equation (9) yields

$$\bar{\mathbb{E}}_t^h[u_t] = (A_u)^h\ u_t + (1+\cdots+A_u^{h-1})A_p(p_t-\theta_{t-1}). \tag{12}$$

Substituting from equation (12) into equation (4) yields:

$$\begin{aligned}\bar{\mathbb{E}}_t^h[\theta_t] &= (A_u)^h\ u_t + (1+\cdots+A_u^{h-1})A_p(p_t-\theta_{t-1}) + \theta_{t-1} \\ &= (A_u)^h\ u_t + \frac{1-A_u^h}{1-A_u}A_p(p_t-\theta_{t-1}) + \theta_{t-1}.\end{aligned} \tag{13}$$

Now, substituting from equation (13) into equation (3) in Proposition 1 yields:

$$\begin{aligned}A_t &= \sum_{h=1}^{\infty}(1-\alpha)\ \alpha^{h-1}\ \left((A_u)^h\ u_t + \frac{1-A_u^h}{1-A_u}A_p(p_t-\theta_{t-1}) + \theta_{t-1}\right) \\ &= \theta_{t-1} + \frac{(1-\alpha)A_u}{1-\alpha A_u}\ u_t + (1-\alpha)\frac{A_p}{1-A_u}\left(\frac{1}{1-\alpha} - \frac{A_u}{1-\alpha A_u}\right)(p_t-\theta_{t-1}) \\ &= \theta_{t-1} + \frac{(1-\alpha)A_u}{1-\alpha A_u}\ u_t + \frac{A_p}{1-\alpha A_u}(p_t-\theta_{t-1}) \\ &= \theta_{t-1} + \frac{(1-\alpha)A_u u_t + A_p(p_t-\theta_{t-1})}{1-\alpha A_u}.\end{aligned} \tag{14}$$

Note that, using (11), if $\sigma_\eta \to \infty$, equation (14) simplifies to equation (7).

For given θ_{t-1} and u_t, aggregate action A_t takes different values for different values of the public signal p_t, depending on the idiosyncratic error term η_t in the public signal. The average public signal, for given θ_{t-1} and u_t, is $E[p_t|u_t] = \theta_{t-1} + u_t$. Then, averaging over the public signal noise, equation (14) becomes:

$$E[A_t|u_t, \theta_{t-1}] = \theta_{t-1} + \phi_p u_t, \text{ where } \phi_p = \frac{(1-\alpha)A_u + A_p}{1-\alpha A_u}.$$

From (11), $\lim_{\sigma_\eta \to 0} \phi_p = 1$ and $\lim_{\sigma_\eta \to \infty} \phi_p = \phi$. Comparing with ϕ in Proposition 2 yields:

$$\phi_p - \phi = \frac{(1-\alpha)A_u + A_p}{1-\alpha A_u} - \frac{(1-\alpha)\beta}{1-\alpha\beta} = \frac{\sigma_\epsilon^4 \sigma_u^2}{(\sigma_\epsilon^2 + (1-\alpha)\sigma_u^2)(\sigma_\epsilon^2\sigma_u^2 + \sigma_\eta^2(\sigma_\epsilon^2 + (1-\alpha)\sigma_u^2))}.$$

Thus, $\phi_p - \phi > 0$ and $\phi_p - \phi$ is increasing in α. Moreover,

$$\frac{d\phi_p}{d\sigma_\eta^2} = -\frac{\sigma_\epsilon^4 \sigma_u^2}{\sigma_\epsilon^2\sigma_u^2 + \sigma_\eta^2(\sigma_\epsilon^2 + (1-\alpha)\sigma_u^2))^2} < 0.$$

Thus, reducing the noise in the public signal (less σ_η^2) raises ϕ_p. □

Proof of Proposition 4: From (2),

$$\begin{aligned}
a_{it} &= (1-\alpha)\mathbb{E}_{it}[\theta_t] + \alpha\ \mathbb{E}_{it}[A_t] \\
&= (1-\alpha)\mathbb{E}_{it}[\theta_{t-1} + u_t] + \alpha\ \mathbb{E}_{it}\left[\theta_{t-1} + \frac{(1-\alpha)A_u u_t + A_p(p_t - \theta_{t-1})}{1-\alpha A_u}\right] \quad \text{(from (14))} \\
&= \theta_{t-1} + (1-\alpha)\mathbb{E}_{it}[u_t] + \alpha\ \frac{(1-\alpha)A_u}{1-\alpha A_u}\mathbb{E}_{it}[u_t] + \alpha \frac{A_p(p_t - \theta_{t-1})}{1-\alpha A_u} \\
&= \theta_{t-1} + \frac{1-\alpha}{1-\alpha A_u}\mathbb{E}_{it}[u_t] + \alpha \frac{A_p(p_t - \theta_{t-1})}{1-\alpha A_u} \\
&= \theta_{t-1} + \frac{1-\alpha}{1-\alpha A_u}(A_u(u_t + \epsilon_{it}) + A_p(p_t - \theta_{t-1})) + \alpha \frac{A_p}{1-\alpha A_u}(p_t - \theta_{t-1}) \quad \text{(from (8) and (10))} \\
&= \theta_{t-1} + \frac{(1-\alpha)A_u}{1-\alpha A_u}u_t + \frac{(1-\alpha)A_u}{1-\alpha A_u}\epsilon_{it} + \frac{A_p}{1-\alpha A_u}(p_t - \theta_{t-1}).
\end{aligned}$$

Thus,

$$\begin{aligned}
a_{it} - \theta_t &= \frac{(1-\alpha)A_u}{1-\alpha A_u}u_t + \frac{(1-\alpha)A_u}{1-\alpha A_u}\epsilon_{it} + \frac{A_p}{1-\alpha A_u}(p_t - \theta_{t-1}) - u_t \\
&= \frac{(A_p + A_u - 1)u_t + A_p\ \eta_t + (1-\alpha)A_u\ \epsilon_{it}}{1-\alpha A_u} \quad \text{(substituting } p_t - \theta_{t-1} = u_t + \eta_t\text{)}.
\end{aligned}$$

Thus,

$$\begin{aligned}(a_{it}-\theta_t)^2 &= \frac{(A_p+A_u-1)^2u_t^2+A_p^2\ \eta_t^2+(1-\alpha)^2A_u^2\ \epsilon_{it}^2}{(1-\alpha A_u)^2} \\ &+ \frac{2(A_p+A_u-1)u_tA_p\ \eta_t+2(A_p+A_u-1)u_t(1-\alpha)A_u\ \epsilon_{it}+2A_p\ \eta_t(1-\alpha)A_u\ \epsilon_{it}}{(1-\alpha A_u)^2}.\end{aligned}$$

Thus,

$$\int(a_{it}-\theta_t)^2di = \frac{(A_p+A_u-1)^2u_t^2+A_p^2\ \eta_t^2+(1-\alpha)^2A_u^2\ \sigma_\epsilon^2+2(A_p+A_u-1)A_p\ u_t\eta_t}{(1-\alpha A_u)^2}.$$

Thus,

$$E[\int(a_{it}-\theta_t)^2di] = \frac{(A_p+A_u-1)^2\sigma_u^2+A_p^2\ \sigma_\eta^2+(1-\alpha)^2A_u^2\ \sigma_\epsilon^2}{(1-\alpha A_u)^2}, \tag{15}$$

where we recognize that if $\alpha=0$, equation (15) simplified to $\frac{\sigma_u^2\sigma_\eta^2\sigma_\epsilon^2}{\sigma_u^2\sigma_\eta^2+\sigma_u^2\sigma_\epsilon^2+\sigma_\epsilon^2\sigma_\eta^2}$, which is the variance of $\theta_t|\theta_{t-1},p_t,x_{it}$. Differentiating with respect to σ_η^2 yields:

$$\frac{dE[\int(a_{it}-\theta_t)^2di]}{d\sigma_\eta^2} = \frac{\sigma_\epsilon^4\sigma_u^4(\sigma_\epsilon^2\sigma_u^2+\sigma_\epsilon^2\sigma_\eta^2+(1-\alpha)(1-2\alpha)\sigma_u^2\sigma_\eta^2)}{(\sigma_\epsilon^2\sigma_u^2+\sigma_\eta^2\sigma_\epsilon^2+(1-\alpha)\sigma_\eta^2\sigma_u^2)^3}. \tag{16}$$

Thus, if $\sigma_\epsilon^2+(1-\alpha)(1-2\alpha)\sigma_u^2\geq 0$ (in particular, if $\alpha\leq 1/2$), the above derivative is strictly positive. If, instead, $\sigma_\epsilon^2+(1-\alpha)(1-2\alpha)\sigma_u^2<0$, the above derivative is strictly positive if and only if σ_η^2 is sufficiently small. □

Proof of Proposition 5: To obtain $E[\int(a_{it}-\theta_t)^2di]$ with only x'_{it}, first let $\sigma_\eta\to\infty$ in (15), and then substitute σ_ϵ^2 with $\sigma_{\epsilon'}^2=\frac{\sigma_\epsilon^2\sigma_\eta^2}{\sigma_\epsilon^2+\sigma_\eta^2}$. Using (11), and recognizing that $\lim_{\sigma_\eta\to\infty}A_p\sigma_\eta^2=0$, the first step yields:

$$\lim_{\sigma_\eta\to\infty}\frac{(\beta-1)^2\sigma_u^2+(1-\alpha)^2\beta^2\ \sigma_\epsilon^2}{(1-\alpha\beta)^2}=\frac{\sigma_\epsilon^2\sigma_u^2(\sigma_\epsilon^2+(1-\alpha)^2\sigma_u^2)}{(\sigma_\epsilon^2+(1-\alpha)\sigma_u^2)^2}.$$

Substituting σ_ϵ^2 with $\sigma_{\epsilon'}^2$ yields:

$$\frac{\sigma_{\epsilon'}^2\sigma_u^2(\sigma_{\epsilon'}^2+(1-\alpha)^2\sigma_u^2)}{(\sigma_{\epsilon'}^2+(1-\alpha)\sigma_u^2)^2}. \tag{17}$$

Now, subtracting (15) from (17) yields:

$$\begin{aligned}&\Delta = E[\int(a_{it}-\theta_t)^2di]_{(x'_{it})}-E[\int(a_{it}-\theta_t)^2di]_{(x_{it},p_t)} \\ &= \frac{\alpha^2\sigma_\eta^4\sigma_\epsilon^4\sigma_u^4}{(\sigma_\eta^2\sigma_\epsilon^2+(1-\alpha)(\sigma_\eta^2+\sigma_\epsilon^2)\sigma_u^2)^2(\sigma_\epsilon^2\sigma_u^2+\sigma_\eta^2(\sigma_\epsilon^2+(1-\alpha)\sigma_u^2))^2}\ (\sigma_\epsilon^4(\sigma_\eta^2+\sigma_u^2)-(1-\alpha)^2(\sigma_\eta^2+\sigma_\epsilon^2)\sigma_u^4).\end{aligned}$$

As expected, $\lim_{\alpha \to 0} \Delta = 0$, because only the amount information matter; and $\lim_{\alpha \to 1} \Delta > 0$, because then citizens put a lot of weight of the pre-existing public information from the previous period, which need to be countered by new public information about θ_t. Moreover, for any $\alpha > 0$, the setting with both public and private signals (x_{it}, p_t) is a normative improvement over the setting with only private signals (x'_{it}) if and only if $\Delta > 0$, i.e., if and only if

$$\left(\frac{\sigma_\epsilon^2}{\sigma_u^2} \right)^2 > (1 - \alpha)^2 \, \frac{\sigma_\eta^2 + \sigma_\epsilon^2}{\sigma_\eta^2 + \sigma_u^2}.$$

The result follows from inspection of this inequality. □