

RESEARCH ARTICLE

Soft Regulation with Crowd Recommendation: Coordinating Self-Interested Agents in Sociotechnical Systems under Imperfect Information

Yu Luo¹, Garud Iyengar², Venkat Venkatasubramanian^{1,2*}

1 Department of Chemical Engineering, Columbia University, New York, NY, United States of America, **2** Department of Industrial Engineering and Operations Research, Columbia University, New York, NY, United States of America

☞ These authors contributed equally to this work.

* venkat@columbia.edu



OPEN ACCESS

Citation: Luo Y, Iyengar G, Venkatasubramanian V (2016) Soft Regulation with Crowd Recommendation: Coordinating Self-Interested Agents in Sociotechnical Systems under Imperfect Information. PLoS ONE 11 (3): e0150343. doi:10.1371/journal.pone.0150343

Editor: Long Wang, Peking University, CHINA

Received: June 8, 2015

Accepted: February 14, 2016

Published: March 15, 2016

Copyright: © 2016 Luo et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data are within the paper.

Funding: This work is supported by Columbia University.

Competing Interests: The authors have declared that no competing interests exist.

Abstract

Regulating emerging industries is challenging, even controversial at times. Under-regulation can result in safety threats to plant personnel, surrounding communities, and the environment. Over-regulation may hinder innovation, progress, and economic growth. Since one typically has limited understanding of, and experience with, the novel technology in practice, it is difficult to accomplish a properly *balanced* regulation. In this work, we propose a control and coordination policy called *soft regulation* that attempts to strike the right balance and create a collective *learning* environment. In soft regulation mechanism, individual agents can accept, reject, or partially accept the regulator's recommendation. This non-intrusive coordination does not interrupt normal operations. The extent to which an agent accepts the recommendation is mediated by a confidence level (from 0 to 100%). Among all possible recommendation methods, we investigate two in particular: the best recommendation wherein the regulator is completely informed and the crowd recommendation wherein the regulator collects the crowd's average and recommends that value. We show by analysis and simulations that soft regulation with crowd recommendation performs well. It converges to optimum, and is as good as the best recommendation for a wide range of confidence levels. This work sheds a new theoretical perspective on the concept of the wisdom of crowds.

Introduction

Regulating emerging technologies is challenging, and often controversial; requiring a careful trade-off between stability, security, performance, and cost in an *uncertain* environment. Recent examples of emerging technologies include hydraulic fracturing, carbon sequestration, deep sea mining, geoengineering, and personalized medicine. Hydraulic fracturing, for

example, has grown to be a transforming force in the petrochemical industries in recent years with its proponents and opponents debating passionately about its benefits and costs to the society with the attendant regulatory challenges [1, 2].

The regulator's (or central planner's) dilemma with regard to emerging technologies is to strike the appropriate *balance* in regulation. Under-regulation can result in damage to plant personnel, surrounding communities, and the environment. Over-regulation, on the other hand, can hamper economic growth and security. When a technology is new, the inherent risks and benefits are not immediately obvious and only become clear over time, making it harder for the regulatory agency to strike the correct balance in the early stages. This uncertainty necessitates a framework that allows for a very close collaboration between the regulatory agency and the regulated entities that have direct access to field performance, and hence have direct knowledge of what worked and what did not.

In a typical regulatory environment involving conventional technologies, regulators issue mandates that have to be followed by the regulated agents. The agents face fines and other punitive consequences for non-compliance. In this paper, we will call this approach *hard regulation*. We argue that hard regulation is not very effective for regulating emerging technologies. Hard regulation also hinders innovation [3]. The regulation of the Internet illustrates these issues very well. Laws like Digital Millennium Copyright Act (DMCA) and Stop Online Piracy Act (SOPA) have been criticized [4–8], as they arguably “reduce freedom of expression and undermine the dynamic, innovative global Internet.” In addition, while attempting to protect intellectual property, these laws hurt computer security by inhibiting research on security related issues [9]. During the period when a new technology is still maturing, the regulator is just as unsure as the regulated agents about the risk-benefit trade-off, and therefore, hard regulation, through its unintended consequences, could potentially do more harm than good. Instead of issuing potentially misdirected mandates, the regulator and the agents should jointly participate in learning about the emerging technology and its payoff structure. The focus of this work is on how to achieve this elusive goal through an intellectual framework that facilitates both *control* and *learning* in sociotechnical systems.

Control and learning are essential elements in managing risk and regulating behavior in sociotechnical systems (see Table 1). In a purely technical setting, i.e., when all the elements of the systems are machines, the common practice to maintain an efficient and stable system is to use *hard* control where the entities follow strictly specified policies. Process control, robotics, etc., are all examples of hard regulation or hard control. When there is no reliable model or a desirable set point available, one needs to simultaneously learn and control the system dynamics. We call techniques, such as machine learning, stochastic approximation, etc., *hard* learning techniques since they also require the entities to follow strict instructions.

However, in a sociotechnical system with active human participants, hard control, or strict mandates, may not always be appropriate. Mandates can potentially do more harm than good as we argued earlier. A more appropriate course of action would be to offer options to agents that are adopted only when they are incentive compatible. We call such approaches *soft* control [10, 11]. Examples of this approach include the soft paternalism approach for modifying social

Table 1. Control and learning in sociotechnical systems.

	Control	Learning
Hard	feedback control, model predictive control, hard regulation, robot formation, laws, etc.	machine learning, stochastic approximation, Kalman filter, evolutionary dynamics, etc.
Soft	persuasion, soft paternalism, peer pressure, social engineering, mechanism design, etc.	social sensing, social learning, pervasive mobile computing, etc.

doi:10.1371/journal.pone.0150343.t001

behavior [12] wherein carefully designed options “nudge” people to make better decisions [13]; or policy teaching [11] wherein the regulator allocates rewards in such a way that the induced action of agents maximizes the regulator’s value. Other examples include efforts by utility companies to induce consumers to minimize power wastage by reporting average consumption [14]; and health tracking devices, e.g., *Fitbit*, *Jawbone Up*, *Nike FuelBand*, etc., that all incorporate social nudging to motivate physical activities. The soft control policy using peer pressure is shown to promote cooperation in these and other settings, both in theory and in practice [15–18].

As in the case with hard control, soft control can be used only when there is a reliable model and a well-defined set point. Soft paternalism and similar social mechanisms are effective because we understand saving energy and staying physically active are the right things to do. What if we do not know what is best for the agents? Soft learning is a class of learning mechanisms that appropriately incentivize agents in a social network to aggregate important information. Examples of soft learning include social sensing and social learning [18–20] in the context of real-time traffic information and online reviews (such as Yelp).

In this work, we propose *soft regulation* as a new regulatory paradigm that combines features of soft control and soft learning. The regulator aggregates key system-level statistics in a privacy preserving [21] manner (individuals do not need to explicitly disclose their states) and shares these statistics with all agents. The agents have the flexibility to accept, reject, or partially accept the recommendations from the regulator based on their own self interests. The recommendations are simply “nudges” [13]. The mechanism does not interrupt the regulated entities who have direct access to field performance. It creates a collective learning environment for both the regulator and the agents. This partial acceptance (or confidence level) of recommendation is a crucial feature of soft regulation. It critically determines the effectiveness of the mechanism. Soft regulation seeks a balance between over- and under-regulation: Agents have the freedom to rely on both individual exploration and social learning.

We expect soft regulation to be effective when the system has the following features:

1. *Imperfect information*: The action-utility payoff structure is poorly understood, i.e., the data are *noisy* and the models are absent or *incomplete*. Each individual may only possess partial information about the unknown process. Agents rely on inaccurate measurements, approximations, or subjective evaluations to optimize. Later we will discuss how noise, paradoxically, is necessary for soft regulation to add value.
2. *Weak interaction*: The agents can optimize their own actions without taking into consideration the response of other agents, i.e., each’s utility or payoff is only a function of the agent’s own state, and the optimal set point is identical among agents. A good example of such a setting is the initial stages of a new technology; the resources being exploited are abundant and the profits of the agents are not limited by competition but by their ability to exploit the resource effectively. Although the reward an agent receives while operating at a set point may vary, the set point itself, however, is likely to be identical or at least restricted to a narrow range. The discovered set points (by soft regulation or traditional methods) will later become the industry standards when the technology matures. Another example of setting with weak or no interaction is when humans improve their own health conditions by changing habits, medications, or even environments. The interaction among agents is usually minimal. Although each has his/her own unique physiological configurations, grouped by characteristics such as age, gender, profession, etc., they are likely to exhibit common optimal set points within groups.
3. *Bounded rationality*: Agents are autonomous and self-interested, and they *always* move in a direction that locally improves utility, subject to available information.

Soft regulation creates a feedback system where agents have the freedom to choose to accept this feedback. Feedback has long been recognized as an essential feature of complex adaptive systems where causes and effects are intertwined. There have been several attempts over the years to understand the dynamics of social systems in terms of feedback control (see, e.g., [22–25]). While such contributions are useful advances, much of this work, however, is conceptual and qualitative. In contrast, soft regulation is a practical and quantitative methodology—self-interested agents use the feedback from past outcomes to determine future actions, and the regulator provides all the agents a feedback that aggregates system-level information. One can extend this to include group, organization, even societal-level feedback loops. In this study, we focus on the bottom two levels of such a feedback control hierarchy.

There are several ways to generate feedbacks for soft regulation. We focus on two methodologies in this paper: the best recommendation and the crowd recommendation. As the name suggests, best recommendation corresponds to the case where the regulator has full information and computes the feedback by solving a *centralized* optimal control problem. The crowd recommendation on the other hand, is simply the average of the agents' actions. We show that, despite its simplicity, crowd recommendation is as good as the best recommendation for a wide range of confidence levels. This is unsurprising. The collective wisdom of groups has been acknowledged in the literature; see, e.g., the Condorcet's jury theorem [26] and popular bestsellers such as *The Wisdom of Crowds* [27] and *Wiser: Getting Beyond Groupthink to Make Groups Smarter* [28]. In this work, we propose a control-theoretic and mathematical framework that goes beyond one-time predictions and investigates the effectiveness of the wisdom of crowds for optimizing a process using continuously refined information.

The scope of this paper is twofold. We first introduce the concept of soft regulation where agents have the freedom to partially accept the feedback from the regulator. We then investigate a unique and practical way of implementing soft regulation—regulator issuing crowd recommendation. In the remaining paper, we formally define the problem and show by analysis and simulation that soft regulation with crowd recommendation indeed performs well.

Background

In this paper, we analyze a stylized model of soft regulation that preserves the essential features discussed in the introduction. The system consists of one regulator and n agents. Agent i wants to select an action x_i that maximizes the value of the real-valued and strongly concave utility function $f_i(x_i)$ over a convex compact set $\mathbb{X} \subseteq \mathbb{R}$. We assume that although the individual utility functions f_i are different for each agent, the optimal set point $\theta^* = \arg \max_{x \in \mathbb{X}} f_i(x)$ is identical across agents.

We assume that the utility function $f_i(x_i)$ is not explicitly known, nor is it deterministic; agents cannot solve the optimization problem explicitly. In theory, by averaging out the noise, one can obtain a more accurate mapping of the utility function. However, in our setting, each sample corresponds to actual utility each agent receives; therefore, they might not have the incentive to oversample at the location where the utility is low. The agents update individual actions using the following dynamics:

$$\bar{x}_i = g_i(x_i)$$

where g_i denotes the optimization algorithm used by the i -th agent. In practice, g_i can be any function that maps an old action x_i to a new action \bar{x}_i . In order to converge to the optimal θ^* , the function must satisfy regularity conditions. More specifically, g_i should converge to a unique fixed point regardless of the initial value of x_i . Although the conclusions of this paper are applicable to a wide range of optimization methods, here we focus on a commonly used

algorithm named Kiefer-Wolfowitz stochastic gradient method [29] where

$$g_i(x_i) = x_i + \frac{a(t)}{c(t)} \cdot (f_i(x_i + c(t)) - f_i(x_i - c(t))). \quad (1)$$

At time t , the i -th agent samples the payoff twice at the vicinity of its current state x_i , which is only known to the agent. The parameters $a(t)$ and $c(t)$ are known and predefined. The agent then computes the next step according to Eq (1). This algorithm is guaranteed to converge in probability when

$$\mathbb{E}[f_i - \mathbb{E}[f_i]]^2 < \infty, \quad \lim_{t \rightarrow \infty} c(t) = 0, \quad \sum_{t=1}^{\infty} a(t) = \infty, \quad \sum_{t=1}^{\infty} \frac{a(t)^2}{c(t)^2} < \infty.$$

We call a setting where an agent updates its action based on its own measurement the *open loop* scenario (or asocial learning as in [20]).

In the soft regulation setting the regulator computes a feedback recommendation u . The agents then combine u with $\bar{x}_i = g_i(x_i)$ to compute a new action x_i^+ in the following manner:

$$x_i^+ = h_i(x_i) \equiv (1 - \beta_i)g_i(x_i) + \beta_i u = (1 - \beta_i)\bar{x}_i + \beta_i u$$

where $\beta_i \in [0, 1]$ or $[0, 100\%]$ is a measure of the *confidence* that the i -th agent puts on the feedback, and is therefore, called the confidence level. Note that we use x_i^+ , $x_{i,t+1}$, and $x_i(t+1)$ notations interchangeably. Subscript t is omitted in most situations. The confidence level β plays an important role in the resulting dynamics. Action changes are relatively independent of recommendation for agents with small β (the *explorers*), and action remains in the vicinity of u for agents with large β (the *followers*).

Note that soft regulation is not an example of direct social learning as described in [20]: There is no “best agent” to follow because the payoffs are private information and noisy. That said, explorers do resemble the asocial innovators and the followers resemble the copying agents in the social learning setting [20]. The confidence level of an agent may be indirectly related to peer pressure [15–18]: The followers experience a higher peer pressure than the explorers, and therefore, set a higher value of β . Also note that $h_i(x_i)$ can be re-written as follows:

$$h_i(x_i) = x_i + (1 - \beta_i)(\bar{x}_i - x_i) + \beta_i(u - x_i).$$

The soft regulation feedback function resembles the feedback seen in bird flocks and swarm intelligence [30].

When the regulator is fully informed about the functions f_i , g_i , and β_i , the optimal feedback u^* can be computed explicitly by solving the following centralized optimal control problem that maximizes social welfare (sum of utilities) over the projected trajectory:

$$\begin{aligned} &\text{maximize} \quad \sum_t w(t) \sum_i f_i(x_i(t)) \\ &\text{subject to} \quad x_i(t+1) = (1 - \beta_i)g_i(x_i(t)) + \beta_i u(t) \end{aligned}$$

where the time varying weight $w(t)$ can favor either the present or future. We call the solution u^* to this problem the best recommendation.

Since the function f_i , g_i , and the parameter β_i are only privately known to the agents, in practice, it is unlikely that the regulator knows the functions and the parameters. Following [14, 27, 28], we assume that the regulator reports the average, i.e., $u = \frac{1}{n} \sum_i x_i$. We call this recommendation the crowd recommendation. Note that using privacy preserving computations [21], the

regulator can compute the crowd recommendation without ever learning any individual input x_i . In the rest of this paper, we demonstrate that the crowd recommendation ensures the convergence to the optimal set point; moreover, it is as good as the best recommendation for a wide range of confidence levels.

Optimality, robustness, and efficiency of soft regulation

Optimality, robustness, and efficiency are three important characteristics of a mechanism. We define an *optimal* process as one that converges to the maximum utility eventually, a *robust* process as one that the equilibrium can be restored when subjected to disturbances, and an *efficient* process as one that converges quickly to the optimum.

We first show that soft regulation with crowd recommendation is *optimal* and *robust* when subjected to bounded noises if the individual confidence level β_i satisfies $0 \leq \beta_i < 100\%$. We describe the essential steps of the proof in this section. Interested readers can find detailed derivations in [S1 Appendix](#).

The nominal (noise free) optimality of soft regulation can be proved using contraction. Recall our assumption that an agent always adopts an optimal algorithm g_i to find θ^* that maximizes utility f_i . In other words, the map defined by g_i converges to θ^* regardless of the initial condition, i.e., g_i is a contraction and θ^* is its unique fixed point. Consequently, there exists a constant $c < 1$ such that $|g_i(x) - g_i(y)| \leq c|x - y|$. If g_i is differentiable; then we have $|g'_i(x)| < 1$ for all $x \in \mathbb{X}$. The soft regulation process is $h_i = (1 - \beta_i)g_i + \beta u$ where $u = \frac{1}{n} \sum_i x_i$ is the crowd recommendation. We define column vector $\mathbf{x} = [x_1, \dots, x_n]^T$ as the system state and $\mathbf{x}^+ = H(\mathbf{x})$ as the system-wide soft regulation map. Whether a map contracts is determined by the eigenvalues of its Jacobian matrix. We can show that the largest absolute eigenvalue of the Jacobian is strictly less than 1 whenever $0 \leq \beta_i < 100\%$. Thus, soft regulation with crowd recommendation is also optimal. Note that this proof is applicable to a completely heterogeneous system where both the optimization algorithm g_i and the confidence level β_i are different among individuals. See [S1 Appendix](#) for detailed derivations.

According to stability theory [31], the optimal state $\mathbf{x}^* = [\theta^*, \dots, \theta^*]^T$ is robust to bounded disturbances if and only if there exists a *continuous* Lyapunov function for the soft regulation process $\mathbf{x}^+ = H(\mathbf{x})$. We show that $V(\mathbf{x}) = \|\mathbf{x} - \mathbf{x}^*\|_1$, i.e., the Manhattan distance or ℓ_1 -norm between the current state \mathbf{x} and the optimal state \mathbf{x}^* , is a Lyapunov function for the process. The contraction result is used to show that such function $V(\mathbf{x})$ is acceptable and the optimal state is robust. This completes the robustness proof.

We use current mean squared error (MSE) to indicate the efficiency of soft regulation, i.e.,

$$\text{MSE} = \frac{1}{n} \|\mathbf{x} - \mathbf{x}^*\|_2^2.$$

Here we are interested in how β affects the MSE. We simplify the analysis by setting $\beta_i \equiv \beta$. From mean value theorem for vector-valued function, we have

$$\mathbb{E}[\text{MSE}^+] \leq m^2 \text{MSE} + (1 - \beta)^2 \sigma_\omega^2$$

where $m = (1 - \beta) \max_{1 \leq i \leq n, x \in \mathbb{X}} |g'_i(x)| + \beta < 1$ is the upper bound of contraction and $\sigma_\omega < \infty$ is the standard deviation of the zero-mean noise imposed on g_i . The MSE evolution exhibits first order dynamics: Increasing β (or increasing m) decreases the steady state expected MSE; however, a larger β leads to a larger m , and therefore, a weaker contraction and slower convergence. In a noisy system where σ_ω is large, it is advisable for the agents to rely more on the recommendation using a larger β . There is a trade-off between accuracy and speed. This further

implies that at a finite time t , there should exist an optimal β^* between 0 and 100% such that the MSE is the smallest.

Simulation of finite-time dynamics

For practical applications, it is more important to understand the transient or *finite-time* dynamics of soft regulation, and more specifically, the role of confidence level in setting the transient performance. We are able to illustrate with a few additional assumptions about the system.

Recall that, we assume the underlying action-utility payoff function $f_i(x_i)$ to be strongly concave. In this section, our analysis will be focused on the *simplest* concave function, namely an identical and quadratic utility function $f_i(x_i) = f(x_i) = -k(x_i - \theta^*)^2 + \omega$. Identical agents are helpful in identifying the effect of confidence level. The results can be readily extended to heterogeneous systems. In order to study the convergence behavior, one can without loss of generality, assume that $\theta^* = 0$. This particular choice for f is motivated by the fact that any strongly concave function can be approximated by a quadratic function near its optimum. The noise is $\omega \sim \mathcal{N}(0, \sigma_\omega)$. Agents only observe the noisy function values—the underlying structure is not known to the agents.

We define the optimization *efficiency* as the percent reduction in MSE:

$$\eta_t \equiv \frac{\text{MSE}_{t_0} - \text{MSE}_t}{\text{MSE}_{t_0}} \times 100.$$

The efficiency is 100 when the system reaches optimum. We simulate the agent dynamics in NetLogo. The agent set is randomized by a fixed random seed in the program to ensure consistency. For each set of parameters, we run the simulation five times and take the average. The observed deviation was insignificant; therefore we omitted error bars. The model parameter values are listed in Table 2. The noise $\omega \sim \mathcal{N}(0, \sigma_\omega)$ has the same variance as another random variable $\epsilon \sim \text{unif}(-\sqrt{3}\sigma_\omega, \sqrt{3}\sigma_\omega)$. We chose σ_ω to be $200/\sqrt{3}$ so that it is computationally equivalent to a uniform ± 200 noise. The parameters do not represent practical meanings. The particular values are chosen such that the results are easily identifiable.

We first run the simulation for soft regulation with best recommendation. Given the quadratic utility, Kiefer-Wolfowitz algorithm, and system-wide confidence level, the regulator can easily compute best recommendation by solving the optimal control problem introduced earlier. One can obtain the nominal (without noise) system dynamics to be

$$\mathbf{x}(t+1) = (1-\beta)(1-4ka_t)\mathbf{x}(t) + \beta \mathbf{1}u(t).$$

Since the stage cost does not penalize input u , the optimal u^* at stage t can be solved as follows

$$u^*(t) = -\left(\frac{(1-\beta)(1-4ka_t)}{\beta}\right) \cdot \frac{1}{n} \sum_i x_i(t).$$

Table 2. Model parameters.

n	σ_ω	θ^*	k	$a(t)$	$c(t)$
1000	$200/\sqrt{3}$	0	100	$1/t$	$1/(t+200)^{1/3}$

doi:10.1371/journal.pone.0150343.t002




Fig 1. Efficiency of soft regulation with best recommendation.

doi:10.1371/journal.pone.0150343.g001

In Fig 1, we plot the efficiency after 200 iterations against different confidence levels. We observe the efficiency increases monotonically as the confidence level increases. This result is not surprising. As the confidence level increases, the regulator has a stronger influence on the agents, therefore, exerting a more efficient control. Even though for each confidence level, the regulator issues the best recommendation, the recommendation is only effective when the agents choose to listen.

In Figs 2 and 3, we plot the efficiency against confidence level for soft regulation with crowd recommendation. The results from Fig 1 are also included as a reference. It is remarkable that soft regulation with crowd recommendation is as good as the one with best recommendation for a wide range of confidence levels (from 0 to 99%). The real advantage of best recommendation only appears when the confidence level is close to 100%. However, to achieve this best recommendation or even hard regulation, the regulator needs information about utility function, optimization algorithm, and the confidence level. This practice, despite being efficient under the setting of complete information, is costly, impractical, and error prone in practical settings. Especially for hard regulation, additional cost of enforcement needs to be considered.

The results in Figs 2 and 3 confirm our previous insights that the confidence level should be set to a large value but not too close to 100%. The open loop system only reaches about 70% optimum. The system performance is more than 90% optimal when the confidence level is 50% (i.e., the agent takes an average between its own optimization result and the recommendation). We also see a sharp decline in performance when confidence level is too close to 100%. Beyond this “cliff,” the agents explore very little and essentially stay where they are.




Fig 2. Efficiency of soft regulation with crowd recommendation.

doi:10.1371/journal.pone.0150343.g002

In Fig 4, we plot the time progressions of efficiency for different confidence levels. When confidence level is low ($\beta = 0$ or 10%), the MSE increases (efficiency declines) before converging. This is caused by large initial step sizes. As confidence level increases, the system begins convergence earlier. As the confidence level further increases, the system shifts from the regime dominated by exploration to the one dominated by conformity, and the recommendation does not have enough time to converge to optimum before agents start conforming.

In order to better understand the connections between the confidence level and the performance that we hypothesized in the previous section and observed in the simulation results, we now attempt to compute a *closed-form* expression for the system state. Recall that x_i denotes the current state of the i -th agent. The updated state x_i^+ implied by the gradient-based update scheme in Eq (1) is given by

$$x_i^+ = (1 - \beta)(1 - 4ka_t)x_i + \beta u + a_t \cdot \frac{1 - \beta}{c_t} \cdot \hat{\omega}_i \quad (2)$$

where $\hat{\omega}_i \sim \mathcal{N}(0, \sqrt{2}\sigma_\omega)$ is the effective noise resulting from computing the discrete approximation to the gradient in Eq (1). Recall that the crowd recommendation is $u = \sum_i x_i/n$. Thus, the updated u^+ of the recommendation is given by

$$u^+ = [1 - 4ka_t(1 - \beta)]u + a_t \cdot \frac{1 - \beta}{c_t} \cdot \frac{1}{n} \sum_i \hat{\omega}_i.$$




Fig 3. Efficiency of soft regulation with crowd recommendation (large confidence levels).

doi:10.1371/journal.pone.0150343.g003

When $n \gg 1$, the expected $u(t)$ can be treated as a deterministic variable. When β is small, $u(t)$ quickly converges to θ^* . Otherwise, $\mathbb{E}[u(t)]$ can be approximated as follows:

$$\mathbb{E}[u(t)] \approx u(t_0) \exp \left[-4k(1 - \beta) \int_{t_0}^t a_\tau d\tau \right]. \quad (3)$$

For the wisdom of crowds, this implies that a group is smart only when the population is large ($n \gg 1$) and agents are not strongly conforming ($\beta \ll 100\%$). Surowiecki's book [27] shares the same insights. Unlike the averaging method in the book, soft regulation is a continuous feedback process. Even though the open loop ($\beta = 0$) system has the fastest converging recommendation, agents cannot make use of it unless they at least partially accept ($\beta > 0$). This paradox suggests some trade-off and balancing between consensus and efficiency.

In addition, as k increases, $u(t)$ approaches θ^* faster, i.e., a more sensitive utility function implies a more reliable recommendation. Unless an agent can estimate the curvature ($\sim k$ for a quadratic function) of the payoff accurately, it is safer to rely on the recommendation when curvature is larger.

From Eqs (2) to (3) it follows that

$$\mathbb{E}[\text{MSE}(t + 1)] \approx \frac{2a_t^2(1 - \beta)^2 \sigma_\omega^2 / c_t^2 + \beta[2(1 - \beta)(1 - 4ka_t) + 1]u(t)^2}{1 - (1 - \beta)^2(1 - 4ka_t)^2}. \quad (4)$$




Fig 4. Efficiency of soft regulation with crowd recommendation over time.

doi:10.1371/journal.pone.0150343.g004

This approximation agrees closely with the simulation (see lines in all simulation result figures). Interested readers can find detailed derivations in [S1 Appendix](#).

From [Eq \(4\)](#), it follows that MSE converges to 0 (soft regulation is optimal). Meanwhile, when confidence level is low, i.e., $\beta \approx 0$ (explorers), the dependence on $u(t)$ vanishes very quickly, and [Eq \(4\)](#) can be simplified as follows:

$$\mathbb{E}[\text{MSE}(t+1)] \approx \frac{2a_t^2\sigma_\omega^2/c_t^2}{1/(1-\beta)^2 - (1-4ka_t)^2}$$

and MSE monotonically *decreases* as β increases. On the other hand, when confidence level is high, i.e., $\beta \approx 100\%$ (followers), the $u^2(t)$ term dominates σ_ω^2 , one can simplify [Eq \(4\)](#) to

$$\mathbb{E}[\text{MSE}(t+1)] \approx \beta \left\{ u(t_0) \exp \left[-4k(1-\beta) \int_{t_0}^t a_\tau d\tau \right] \right\}^2$$

and MSE monotonically *increases* as β increases. This estimation agrees well with our previous hypotheses and simulation results (see [Figs 2 to 4](#)).

An interesting fact arises from this approximation, i.e., imperfect information is necessary for soft regulation to add value. If the system has very low noise or noise-free, the σ_ω^2 term will be dominated by $u^2(t)$, and an increase in β hurts performance. That is to say, *for a deterministic process, soft regulation with crowd recommendation may not be a good mechanism for agents to adopt.*

In practice, each individual may have a distinct confidence level and personal traits. Modeling such rich details as well as formulating related best recommendation is beyond the scope of this work. Nevertheless, for purposes of illustration, we propose the following adaptive confidence mechanism:

$$\beta(x_i, u) = e^{-b(x_i - u)^2} \quad (b > 0).$$

The rationale for this update scheme is as follows. When an agent's action is far away from the recommendation, the agent is fairly skeptical. Suppose, by incorporating the recommendation, the agent's action moves further away from the recommendation, the agent would rely even less on the regulator. However, when the action comes closer to the recommendation, agent is likely to be more confident about the regulator, and incorporate the recommendation in future updates. One flaw in this adaptive mechanism is that if everyone performs the same action in the beginning, this results in an identical confidence level $\beta = 100\%$ for all agents, and the system will not move at all. This situation might be remedied with an occasional, small perturbation. In Fig 5, we plot two new simulation results, i.e., 1) agents have uniformly distributed (*Dist.*) confidence levels, and 2) agents have uniformly distributed initial confidence levels and the confidence is adaptive (*Dist.+adap.*) according to the update scheme above. We also include previous results with fixed and identical confidence level to the graph. We observe a fairly good performance.




Fig 5. Efficiency of soft regulation with crowd recommendation over time (distributed agents).

doi:10.1371/journal.pone.0150343.g005

Conclusion

We propose a soft regulation framework for coordinating multi-agent systems. In this setting, the regulator's role is to help agents learn, understand, and optimize an unknown process without interrupting normal operations. The essence of this mechanism is to take guided decisions by updating actions using the map $x_i^+ = (1 - \beta_i)\bar{x}_i + \beta_i u$, where \bar{x}_i is the i -th agent's own optimized decision and u is the regulator's crowd recommendation. Self-interested agents have the freedom to choose to partially accept the regulator's recommendation. Soft regulation provides a more balanced coordination: Unlike hard regulation, it does not force the agents; this creates a collective learning environment and avoids possibly erroneous mandates. On the other hand, a soft regulatory system is not under-regulated or uncoupled. The exploration of some agents benefits others. Useful information is shared indirectly instead of being wasted in an asocial learning environment.

We notice the efficiency of soft regulation is impacted by the following factors:

1. *Population*: Because of noise, recommendation is subject to uncertainty. However, when $n \gg 1$, the variance becomes negligible, and the recommendation becomes deterministic (very close to mean) and accurate. This dependence on population size is intuitive: The information aggregated from a large population should be more useful compared to the one from a small population.
2. *Process*: We have proved that soft regulation with crowd recommendation preserves optimality. The advantage of the mechanism, however, is especially pronounced when the system is very noisy and the payoff function is very sensitive. A rule of thumb for the agents would be when a large sensitivity of the process is observed (either because of high noise or large curvature), the agents may be better off relying on the recommendation. Uncertainty drives the system towards cooperation. Soft regulation can potentially stabilize an open loop unstable process. This result also provides some insights on the wisdom of crowds. For example, the average performance can outperform the best individual when the system is very uncertain. In that sense, the "expert" knowledge may not be as useful in an emerging industry as the collective wisdom.
3. *Agent confidence level*: From the mathematical proofs and simulation evidence, we conclude that the best confidence level should be large but not too close to 100%. This is especially true when the system is very noisy and the process is very sensitive. In such setting, agents should put a substantial amount of trust on the regulator's recommendations. Because of the trade-off between consensus and efficiency, in the early stage of soft regulation, the confidence level should be kept low for recommendation to quickly converge. As time proceeds, agents can be more and more confident regarding the recommendation.

Despite the name, soft regulation has applications beyond industrial regulation. The soft regulator module, i.e., $x_i^+ = (1 - \beta_i)\bar{x}_i + \beta_i u$, can be integrated in different control systems and problem-solving scenarios (see Table 3). We only analyze a specific and stylized model in

Table 3. Soft regulation applications.

Action	→	Utility
health behavior: e.g., sleep habit, exercise frequency, diet, etc.	→	health condition e.g., sleep quality, BMI, etc.
operating condition: e.g., T , P , feed ratio, flow rate, catalyst, etc.	→	yield e.g., production rate, etc.
workplace environment: e.g., indoor temperature, lighting, etc.	→	productivity e.g., profitability, etc.
infrastructure planning: e.g., traffic light control, hospital resource, budget allocation, etc.	→	efficiency e.g., congestion time, etc.

doi:10.1371/journal.pone.0150343.t003

this paper to illustrate the efficacy of the mechanism. In practice, soft regulation should be implemented and modified in a case by case manner. For example, when the regulator can obtain more information other than the average action, it is entirely reasonable to formulate a better recommendation based on the richer information set such as trends, histograms, etc. The agents, instead of adjusting confidence level via the method proposed in this paper, can also explore and compare utilities (on a much slower timescale) to adapt new confidence levels. For a large population where centralized information collection is impractical, soft regulation might be plausible on a peer-to-peer basis. All these possibilities will be explored and analyzed in future work.

The medical domain is another applicable area of soft regulation. Powered by mobile phones and wearables, researchers can now collect timely mass medical data (via Apple's ResearchKit [32] for example). Soft regulation is suitable in this scenario because medical research satisfies all three features, i.e., imperfect information (unknown relationships between patient behaviors and health conditions), weak interaction (one patient's condition is not affected by another's), and bounded rationality (patients always wish to improve their own health, however, have limited information). In addition, thanks to the convenience of mobile devices, we expect good participation rate. A large population size further ensures the accuracy of recommendation. Patients can optimize their own health while contributing to medical research. Even if patients do not want to optimize themselves, medical researchers may implement the soft regulation module to do that based on data collected locally. The confidence level can also be explicitly controlled by the service provider. Soft regulation in this setting becomes a crowdsourcing framework. The results in this paper are expected to hold.

This work also has some implications other than our central arguments on control and regulation. It reinforces the idea that an averaged opinion can accurately predict under uncertainty, i.e., the wisdom of crowds, given the population is large, independent, and relevant. Unlike conventional takes on the wisdom of crowds, soft regulation does not stop at collecting average information but also feeds it back to the system. This dynamical mechanism suggests more flexible scenarios and applications.

Supporting Information

S1 Appendix. Proofs and detailed derivations. In the Appendix, we show that soft regulation with crowd recommendation is optimal and robust when subjected to bounded noises when $0 \leq \beta_i < 100\%$. We also discuss how confidence level affects the efficiency of this mechanism. In addition, we formulate a closed-form analytical solution for the simulation case study. (PDF)

Acknowledgments

This work is supported in part by Columbia University.

Author Contributions

Conceived and designed the experiments: VV. Performed the experiments: YL. Analyzed the data: YL GI VV. Wrote the paper: YL GI VV.

References

1. Gold R. The Boom: How Fracking Ignited the American Energy Revolution and Changed the World. Simon and Schuster; 2014.
2. Bloomberg, M, Krupp, F. type; 2014. Available from: <http://nyti.ms/1nEASw5>.

3. Krupp, F. type; 2008. Available from: <http://www.strategy-business.com/article/08201?pg=all>.
4. Congress U. Digital Millennium Copyright Act. Public Law. 1998;(105-304):112.
5. Tribe LH. The "Stop Online Piracy Act" (SOPA) violates the first amendment; 2011.
6. Smith, L. Stop online piracy act. US Government. 2011;.
7. Rai AK. Intellectual property rights in biotechnology: Addressing new technology. *Wake Forest L Rev*. 1999; 34:827.
8. Phillips, M. type; 2012. Available from: <http://www.whitehouse.gov/blog/2012/01/14/obama-administration-responds-we-people-petitions-sopa-and-online-piracy>.
9. Higgins, P. type; 2014. Available from: <http://gizmodo.com/how-drm-harms-our-computer-security-1572826413>.
10. Han J, Li M, Guo L. Soft control on collective behavior of a group of autonomous agents by a skill agent. *Journal of Systems Science and Complexity*. 2006; 19(1):54–62. doi: [10.1007/s11424-006-0054-z](https://doi.org/10.1007/s11424-006-0054-z)
11. Zhang H, Parkes DC. Value-Based Policy Teaching with Active Indirect Elicitation. In: *AAAI*. vol. 8; 2008. p. 208–214.
12. Thaler RH, Sunstein CR. Libertarian paternalism. *American Economic Review*. 2003; p. 175–179. doi: [10.1257/000282803321947001](https://doi.org/10.1257/000282803321947001)
13. Thaler RH, Sunstein CR. *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press; 2008.
14. Brotman, B. type; 2014. Available from: http://articles.chicagotribune.com/2014-03-03/features/ct-energy-comparisons-brotman-talk-0303-20140303_1_energy-hog-energy-efficiency-comed.
15. Kandel E, Lazear EP. Peer pressure and partnerships. *Journal of political Economy*. 1992; p. 801–817. doi: [10.1086/261840](https://doi.org/10.1086/261840)
16. Aharony N, Pan W, Ip C, Khayal I, Pentland A. Social fMRI: Investigating and shaping social mechanisms in the real world. *Pervasive and Mobile Computing*. 2011; 7(6):643–659. doi: [10.1016/j.pmcj.2011.09.004](https://doi.org/10.1016/j.pmcj.2011.09.004)
17. Mani A, Rahwan I, Pentland A. Inducing peer pressure to promote cooperation. *Scientific reports*. 2013; 3. doi: [10.1038/srep01735](https://doi.org/10.1038/srep01735)
18. Shmueli E, Singh VK, Lepri B, Pentland A. Sensing, Understanding, and Shaping Social Behavior. *Computational Social Systems*, *IEEE Transactions on*. 2014;. doi: [10.1109/TCSS.2014.2307438](https://doi.org/10.1109/TCSS.2014.2307438)
19. Krishnamurthy V, Poor HV. A Tutorial on Interactive Sensing in Social Networks. *Computational Social Systems*, *IEEE Transactions on*. 2014;. doi: [10.1109/TCSS.2014.2307452](https://doi.org/10.1109/TCSS.2014.2307452)
20. Rendell L, Boyd R, Cownden D, Enquist M, Eriksson K, Feldman MW, et al. Why copy others? Insights from the social learning strategies tournament. *Science*. 2010; 328(5975):208–213. doi: [10.1126/science.1184719](https://doi.org/10.1126/science.1184719) PMID: [20378813](https://pubmed.ncbi.nlm.nih.gov/20378813/)
21. Abbe EA, Khandani AE, Lo AW. Privacy-preserving methods for sharing financial risk exposures. *The American Economic Review*. 2012; 102(3):65–70. doi: [10.1257/aer.102.3.65](https://doi.org/10.1257/aer.102.3.65)
22. Carver CS, Scheier MF. Control theory: A useful conceptual framework for personality—social, clinical, and health psychology. *Psychological bulletin*. 1982; 92(1):111. doi: [10.1037/0033-2909.92.1.111](https://doi.org/10.1037/0033-2909.92.1.111) PMID: [7134324](https://pubmed.ncbi.nlm.nih.gov/7134324/)
23. Leveson N. *Engineering a safer world: Systems thinking applied to safety*. Mit Press; 2011.
24. Trochim WM, Cabrera DA, Milstein B, Gallagher RS, Leischow SJ. Practical challenges of systems thinking and modeling in public health. *American Journal of Public Health*. 2006; 96(3):538. doi: [10.2105/AJPH.2005.066001](https://doi.org/10.2105/AJPH.2005.066001) PMID: [16449581](https://pubmed.ncbi.nlm.nih.gov/16449581/)
25. Powers WT. Feedback: beyond behaviorism stimulus-response laws are wholly predictable within a control-system model of behavioral organization. *Science*. 1973; 179(4071):351–356. doi: [10.1126/science.179.4071.351](https://doi.org/10.1126/science.179.4071.351) PMID: [4682961](https://pubmed.ncbi.nlm.nih.gov/4682961/)
26. De Condorcet N. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Cambridge University Press; 2014.
27. Surowiecki J. *The wisdom of crowds*. Anchor; 2005.
28. Sunstein C, Hastie R. *Wiser: Getting Beyond Groupthink to Make Groups Smarter*. Harvard Business Review Press; 2014.
29. Kiefer J, Wolfowitz J, et al. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*. 1952; 23(3):462–466. doi: [10.1214/aoms/1177729392](https://doi.org/10.1214/aoms/1177729392)
30. Kennedy J. Particle swarm optimization. In: *Encyclopedia of Machine Learning*. Springer; 2010. p. 760–766.

31. Teel AR. Discrete Time Receding Horizon Optimal Control: Is the Stability Robust? In: Optimal Control, Stabilization and Nonsmooth Analysis. Springer; 2004. p. 3–27.
32. Apple. ResearchKit for Developers; 2015. Available from: <https://developer.apple.com/researchkit/>.