A One-Third Advice Rule Based on a Control-Theoretic Opinion Dynamics Model

Yu Luo, Garud Iyengar, and Venkat Venkatasubramanian

***Abstract*— We commonly seek advice in making decisions. However, multiple empirical studies report that, on average, we shift our own initial decision by only 30% toward external advice after advice is provided. This "egocentric advice discounting" is particularly counterintuitive because we do care a lot about the opinion of our peers. There is significant literature that attempts to explain the egocentric advice discounting and factors that influence this phenomenon; however, this literature is unable to explain why the numerical value of 30% is robust across a number of experimental settings. In this paper, we employ a control-theoretic opinion dynamics model to show that the one-third advice rule—adjusting one's decision about 33.3% toward advice—is in fact distributionally robust for a crowd of decision-makers whose decisions also serve as advice for others. Our results imply that the observed egocentric advice discounting might not be a coincidence; instead, when an individual is faced with insufficient information, the distributionally robust optimal decision is to combine one-third of advice with two-thirds of his/her initial decision. Our theory also suggests that knowing the dispersion of decisions can further help decision-makers optimize advice taking.**

***Index Terms*— Advice taking, control theory, decision-making, judge–advisor system, opinion dynamics, social influence, wisdom of crowds.**

I. Introduction

We often make decisions after polling our peers. This behavior has become even more prevalent in the Internet age. We read reviews and ratings of other consumers before purchasing any product online, visiting a restaurant, or choosing an accommodation. Given the effort put into collecting the ratings and reviews, it is clear that the online retailers believe that these ratings, comments, and popularity of a product strongly influence purchasing decisions. Advice taking—the process of revising one's decision or opinion after receiving advice—is an extensively studied subject [1], [2]. In a previous study with human subjects, we discovered that, on average, people shifted their decisions by some 30% toward external advice when advice was present [3]. After fitting our decision-making model to the experiment data, we estimated that the 30% shift was also optimal in minimizing the cumulative squared decision error.

Many empirical findings on how much we are willing to change our decision based on advice also report similar shifts [1], [4]. In [5], to improve judgment, research participants shifted their judgment by 20%–30% toward advice when advice was provided to them. In [6], subjects revised their initial forecast by shifting it about 33% toward a statistical forecast in the light of this additional piece of information. In [7], subjects placed a weight of 71% on their own estimates (i.e., shifting about 29% toward advice) when advice was given. In a paper presented at the annual meeting of the Society for Judgment and Decision Making in 1999 [8], Soll and Larrick stated that final judgments can be quantified as 80% own initial judgment and 20% peer (i.e., the "80/20" rule).

This phenomenon is particularly counterintuitive because we clearly care about the opinion of our peers. The Roman Emperor Marcus Aurelius famously observed [9] the following.

> It never ceases to amaze me. We all love ourselves more than other people, but care more about their opinion than our own.

The "wisdom of crowds" phenomenon [10]–[12], which explains the efficiency of the market, online review systems, crowdsourced platforms, and so on, relies on the fact that simple average of independent opinions of error-prone decision-makers is often significantly superior to any given opinion. In many circumstances, the consensus decision is significantly superior even if the individuals are erroneous. Consequently, decision-makers should assign at least as much weight to the consensus opinion as to their own. In situations with experts who possess superior information, decision-makers should consider advice even more, not less. Thus, one would expect that the weight on the advice to be greater than 50% with a dispersion over 50%–100% to account for population heterogeneity. The empirical finding of 30% weight on external advice, therefore, needs careful consideration and explanation.

Psychologists explain this "egocentric advice discounting" [13] using internal justifications and self-anchoring [2]. The approach here is to treat the egocentric advice discounting as a behavioral bias that induces decision-makers to discount advice and rely heavily on their own opinions. *But what if*

Manuscript received December 8, 2018; revised February 26, 2019; accepted April 1, 2019. Date of publication May 1, 2019; date of current version June 10, 2019. This work was supported by the Center for the Management of Systemic Risk, Columbia University. *(Corresponding author: Venkat Venkatasubramanian.)*

Y. Luo was with the Department of Chemical Engineering, Columbia University, New York, NY 10027 USA. He is now with the Department of Chemical and Biomolecular Engineering, University of Delaware, Newark, DE 19716 USA.

G. Iyengar is with the Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027 USA.

V. Venkatasubramanian is with the Department of Chemical Engineering, Columbia University, New York, NY 10027 USA (e-mail: venkat@columbia.edu).

Digital Object Identifier 10.1109/TCSS.2019.2910261

the 30% *rule is not a bias, rather a meaningful heuristic for incorporating advice?* In this paper, we attempt to identify minimal constructs that justify such observed phenomenon.

Our motivation is as follows. A quantitative analysis could potentially provide a more precise explanation of the numerical value of 30%. It might reveal whether the egocentric advice discounting is a psychological bias or, in fact, the 30% rule is, perhaps, the result of optimization—have we, the decision-makers, found a robust heuristic over repeated trials?

In what follows, we will argue that the 30% rule is approximately "optimal," given the available information about the uncertain parameters. We assume that the available information is summarized by a distribution. Such an assumption is standard in the decision sciences and risk management literature [14]–[16]. We assume that the uncertain parameters are distributed according to the maximum entropy distribution subject to the available information. Such a distribution is, in a sense, maximally random [17], [18].

When we began analyzing the 30% rule, we found it a curious value—too "rational" as compared to other universal constants, such as π, e, golden ratio, and speed of light and too much of an artifact of a decimal system derived from the fact that humans have ten fingers, having nothing to do with decision-making or advice taking. We noticed that one of the studies that report the 30% shift states the following [6].

> It shows that in general, the final forecast was approximately made by a $2/3 \times$ initial forecast $+$ $1/3 \times$ statistical forecast model—that is, a twice-greater weight for their initial forecast than the statistical forecast.

Even though $1/3$ might be as arbitrary as 30%, it appeared more plausible to us. In the remainder of this paper, we show that $1/3$ emerges as a distributionally robust design in decision-making.

We borrow the term "one-third rule" or "rule of thirds" from photography where important compositional elements should be placed along the thirds, instead of the more intuitive center position. The same counterintuition also applies to taking advice according to our theory.

II. Methods

A model for the decision-making process is crucial in determining the optimal advice taking. Such a model should address two critical characteristics of the process. First, the model should be able to account for how individuals revise their decisions as a response to *asocial* feedback (outcomes of decisions) and *social* feedback (decisions made by others). Feedback plays an important role in this paper, for that we attempt to determine the optimal response to the feedback that would benefit individual decision-makers.

Second, decision-making is usually a continuous process, where decisions evolve over time. Therefore, we require the model to be able to describe both one-off and changing decisions. The Delphi method, for instance, is a systematic and iterative process of generating consensus among experts [19], [20]. Individuals can revise their opinions after reviewing others' opinions (opinions are collected anonymously). A consensus can usually be reached in at least two and no more than ten iterations [21]. Other examples include polls, online reviews, and stock market prices with dynamically updated sources of information. Participants periodically make decisions, receive feedback, and update their decisions.

Opinion dynamics is an active research area in the social sciences [22]–[28], and there are numerous models that describe how opinions evolve and interact [29]. A recurring theme in opinion dynamics is the convergence or divergence of the opinions for a group of individuals (or agents) who have access to other individuals' opinions and can revise their own accordingly. Social feedback is the driving force of opinion dynamics in most models. In this paper, we are focused on a special case of opinion dynamics, in which asocial feedback and social feedback are equally important in shaping decisions. Decisions result in consequences, and these consequences lead to the decision-makers to update and improve their decisions, even in the absence of any social feedback. For example, a purchase decision that results in a product that does not meet a consumer's expectation is likely to discourage the consumer from purchasing the same product in the future. Consequently, in order to understand the benefit of social feedback in improving decisions, one needs to model not only the social influence on consensus forming but also the evolution of decisions when there is no social feedback. This tension between asocial feedback and social feedback, between exploration and exploitation, and between innovation and optimization differentiates the problem structure that we discuss in this paper from the classical opinion dynamics models.

Control theory is a well-established canonical framework that studies dynamical systems with feedback [30]–[33]. Attempts have been made in the past to incorporate feedback control into the studies of social systems [34]–[40], including our previous works [3], [41]. "Perceptual control theory" proposed by an independent psychologist William Powers [37]–[40], for example, explicitly incorporates control theory concepts in psychology and ambitiously aims at becoming the ultimate quantitative model of behavior. Our goal, on the other hand, is much more focused—we want to employ the control theory formalism to understand how individuals revise their decisions when advice is present. We can use control theory and its repertoire of tools to generate predictions and explain the egocentric advice discounting phenomenon quantitatively.

Fig. 1 displays the block diagram for the decision-making process of an individual when there is only asocial feedback present. We illustrate the feedback control on the following health care example. Suppose a diabetic patient sets the insulin dosage to a level θ. This decision interacts with a noisy environment to produce a certain blood sugar level—the decision outcome. This patient then uses a noisy measurement of the blood sugar level to adapt the insulin level to θ^+, mimicking the blood sugar control process of a healthy pancreas.

There is an important difference between the system in Fig. 1 and most opinion dynamics models. In the classical opinion dynamics setting, when social feedback is absent, an agent's opinion: 1) remains constant (e.g., the DeGroot model [22] and the Deffuant–Weisbuch model [42]); 2) is

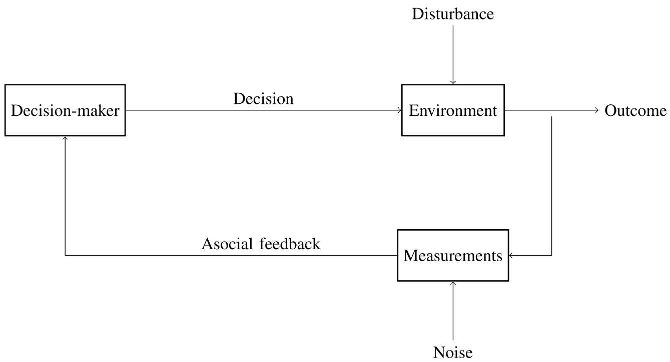

Fig. 1. Feedback loop of individual decision-making process.

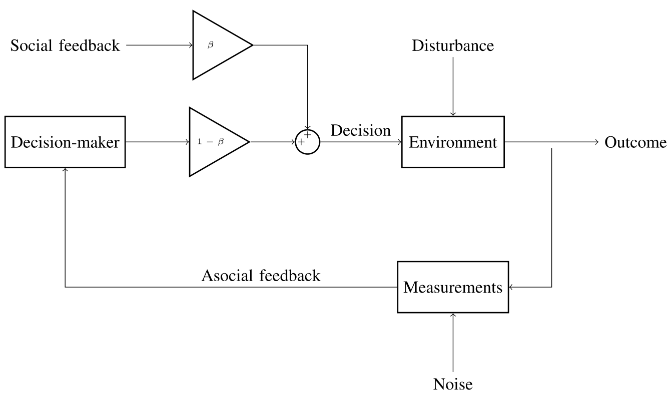

Fig. 2. Feedback loop of individual and social decision-making process.

determined by exogenous variables (e.g., the Friedkin–Johnsen model [23]); or 3) is gradually pulled toward some "truth" (e.g., the Hegselmann–Krause model [43]). Specifically, Hegselmann and Krause [43] emphasize that the agents are not intentionally following such mechanism because they would immediately choose the "truth" as the definite opinion. While the Hegselmann–Krause model is very close to what Fig. 1 describes, our formulation focuses on how decisions evolve based on the consequences of past decisions, instead of an artificial attraction toward the "truth."

In Fig. 2, we display a control-theoretic block diagram for advice taking (when a social feedback is present). In this setting, the updated decision is the weighted sum of the individual's own decision and the social advice. β is the weight on advice. We called β the "degree of social influence" in our previous work [3] because it represents the influence of other people's opinions. In psychology literature (including the opinion dynamics), the terms "advice taking" [5], [44], "weight of advice" [4], [45], and "cautiousness parameter" (in the Deffuant–Weisbuch model [42]) are also used.

If we denote x as the (error of) current decision (i.e., we subtract the optimal decision θ^* from the face value of decision θ and the optimal decision is $x^* = 0$) and x^+ as the updated decision, we can describe an individual's decision-making process as follows:

$$x^+ = g(x) + \omega \tag{1}$$

where the innovation function $g(\cdot)$ maps an old decision to a new one and the process is disturbed by a zero-mean random variable ω. The innovation function describes the decision-making process when a social feedback is absent and new decisions are made based on past decisions and decision outcomes. Under mild regularity conditions, we assume that the optimal decision $x^* = 0$ is a fixed point of the mapping g, i.e., $x^* = g(x^*)$. The decision-making dynamics without advice (see Fig. 1) can also be written as

$$x^+ = \lambda x + \omega \tag{2}$$

where ω is again a zero mean variance σ^2 random variable that models the combined effect of the disturbance and noise. The parameter $\lambda \equiv g(x)/x$ is the ratio between the updated and the old decisions. $|\lambda| < 1$ would imply that, on average, the error of decision decreases after an update. Both λ and σ are *time-varying* parameters that represent the evolving nature of decision-making.

In Fig. 2, the decision update with advice (social feedback) is given by

$$x^+ = (1-\beta)(\lambda x + \omega) + \beta u \tag{3}$$

where the weight β ranges from 0 (not taking any advice) to 1 or 100% (completely following/copying advice) and u is the average decision of all participating agents

$$u \equiv \frac{1}{n}\sum_{i=1}^{n} x_i \tag{4}$$

where i denotes the ith agent in an n-member population ($i = 1, \ldots, n$). The average opinion is particularly attractive because it can be computed in a privacy-preserving manner, i.e., without the individual agents disclosing their true decisions [46].

Assumption 1 (Individual Efficiency): The update of the individual decision in (2) is efficient on average, that is

$$\mathbb{E}_\omega[x^{+2}] < x^2. \tag{5}$$

It is reasonable to assume that when individual errors are large (e.g., at the early stage of optimization), as the individual gathers more information and feedback, naturally his/her new decision should improve (probabilistically). This assumption could, however, be violated when decisions are very close to the optimum. The one-third advice rule is, therefore, only applicable to situations where the optimal decisions (or "truth" [43]) have yet to be identified or accepted by the population.

From (2), we have

$$\mathbb{E}_\omega[x^{+2}] = \lambda^2 x^2 + \sigma^2. \tag{6}$$

In order to satisfy Assumption 1, we have

$$\lambda^2 + \gamma^2 < 1 \tag{7}$$

where $\gamma \equiv \sigma/|x|$ is the scaled noise parameter. Recall that the parameters (λ, σ) in (2) for each agent are possibly random time-varying parameters.

Our objective in this paper is to understand how β affects the squared error of decision $\mathbb{E}_\omega[x^{+2}]$. Let F denote the joint distribution of the set of parameters $\{(\lambda_i, \gamma_i) : i = 1, \ldots, n\}$ that describe the dynamics of all n agents.

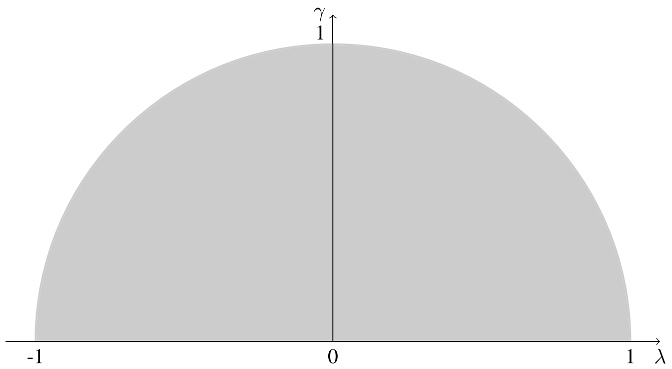

Fig. 3. Semi-unit-circle of λ and γ values that satisfy Assumption 2 (shaded).

Note that in the social feedback setting, the feedback u couples the dynamics of all agents. We are interested in solving the optimization problem from the perspective of an individual decision-maker

$$\beta^* \equiv \arg\min_{\beta} \mathbb{E}_F[x^{+2}(\beta)]. \tag{8}$$

Following the maximum entropy principle [17], [18], we assume that the nature selects the parameters (λ_i, γ_i) in a maximally random manner subject to the constraint $\lambda_i^2 + \gamma_i^2 < 1$ imposed by Assumption 1.

Assumption 2 (Distribution F): The set of parameters (λ_i, γ_i), $i = 1, \ldots, n$, are drawn independently and identically according to the uniform distribution on the semicircle $S = \{(\lambda, \gamma) : \gamma > 0, \lambda^2 + \gamma^2 < 1\}$ (see Fig. 3).

The optimization problem (8) is hard to solve in complete generality. Each individual needs to know others' decisions and dynamics (λ, γ, x, u, and so on). More importantly, each decision-maker needs to know how the advice influences the other decision-makers. Suppose an agent assumes that all other agents completely ignore the social feedback. Then, the choice of β will have no impact on the future evolution of the decisions; thus, the advice will be of little value, and the decision dynamics will reduce to the case without any feedback. A more reasonable assumption about the other agents is that they are also solving a similar problem, based on the fact that individuals are statistically indistinguishable (while the realization of parameters λ and γ is distinct for each unique individual). One may also postulate more refined models of the beliefs that a given decision-maker has about other decision-makers. However, recall that in this paper, we are solving an inverse problem, i.e., to identify the minimal constructs that explain an observed phenomenon. To this end, we begin with the following simple belief structure.

Assumption 3 (Identical Optimization Problem): All individuals assume that all other agents solve the same optimization problem (8).

Hence, solving (8) for each individual under Assumption 3 is equivalent to solving the following program for the population as whole:

$$\beta^* \equiv \arg\min_{\beta} \mathbb{E}_F\left[\mathrm{MSE}^+(\beta)\right] \tag{9}$$

where the mean squared error (MSE) is defined as

$$\mathrm{MSE} \equiv \frac{1}{n}\sum_{i=1}^{n} {x_i}^2. \tag{10}$$

In Section III, we show the optimal $\beta^* = 1/3$. Thus, we have identified a plausible framework for explaining the one-third rule—the agents are risk-averse over the realization of the parameters and assume that all other agents are solving similar optimization problems.

III. Results and Discussion

Since ω_i is a zero mean random variable with variance σ_i^2, it follows that:

$$\begin{aligned}\mathbb{E}_{\omega_i}[\mathrm{MSE}^+] &\equiv \mathbb{E}_{\omega_i}\left[\frac{1}{n}\sum_{i=1}^{n}[(1-\beta)(\lambda_i x_i + \omega_i) + \beta u]^2\right] \\ &= \frac{1}{n}\sum_{i=1}^{n}(1-\beta)^2\left(\lambda_i^2 + \gamma_i^2\right)x_i^2 \\ &\quad + 2\beta(1-\beta)u\frac{1}{n}\sum_{i=1}^{n}\lambda_i x_i + \beta^2 u^2.\end{aligned} \tag{11}$$

From Assumption 2, it follows that:

$$\begin{aligned}\mathbb{E}_{\lambda,\gamma}[\lambda^2 + \gamma^2] &= \frac{1}{2} \\ \mathbb{E}_{\lambda,\gamma}[\lambda x] &= 0.\end{aligned} \tag{12}$$

Taking the expectation of (11), we have

$$\begin{aligned}\mathbb{E}_F[\mathrm{MSE}^+] &= \mathbb{E}_{\lambda_i,\gamma_i}\Bigg[\frac{1}{n}\sum_{i=1}^{n}(1-\beta)^2\left(\lambda_i^2 + \gamma_i^2\right)x_i^2 \\ &\quad + 2\beta(1-\beta)u\frac{1}{n}\sum_{i=1}^{n}\lambda_i x_i + \beta^2 u^2\Bigg] \\ &= \frac{1}{2}(1-\beta)^2\mathrm{MSE} + \beta^2 u^2.\end{aligned} \tag{13}$$

Thus, the optimization problem (9) reduces to

$$\min_{\beta}\left\{\frac{1}{2}(1-\beta)^2\mathrm{MSE} + \beta^2 u^2\right\}. \tag{14}$$

The presence of the feedback u in the objective couples the decisions β across time, resulting in a dynamic optimization. Furthermore, the solution of this dynamic problem is quite sensitive to details of the parameters. Since our goal here is to posit a simple rule, we take a minimax approach and posit the agents choose β by solving the minimax problem

$$\begin{aligned}\beta^* &\equiv \arg\min_{\beta}\max_{u}\mathbb{E}_F[\mathrm{MSE}^+] \\ &= \arg\min_{\beta}\left[\frac{1}{2}(1-\beta)^2\mathrm{MSE} + \beta^2\max_{u} u^2\right].\end{aligned} \tag{15}$$

Since $f(x) = x^2$ is a convex function of x, Jensen's inequality [47] implies that f of the average of $x_1, \ldots, x_n$ is less than or equal to the average of $f(x_1), \ldots, f(x_n)$

$$u^2 \equiv \left(\frac{1}{n}\sum_{i=1}^{n}x_i\right)^2 \le \frac{1}{n}\sum_{i=1}^{n}x_i^2 \equiv \mathrm{MSE}. \tag{16}$$

As a result

$$\begin{aligned}\beta^* &= \arg\min_{\beta}\left[\frac{1}{2}(1-\beta)^2\mathrm{MSE} + \beta^2 \max_u u^2\right]\\ &= \arg\min_{\beta}\left[\frac{1}{2}(1-\beta)^2 + \beta^2\right]\mathrm{MSE}\\ &= \frac{1}{3}.\end{aligned} \tag{17}$$

This distributionally robust solution $\beta^* = 1/3$ corresponds to the worst case distribution, where $x_i \equiv u$, i.e., there is no diversity in the population. In such a situation, it is wise to be conservative in taking the population advice. Next, we look at how additional information about the system, such as knowing the dispersion of decisions, influences the optimal β^*.

The coefficient of variation is defined as $\mathrm{CV} \equiv \sigma_x/|u|$, where σ_x is the sample standard deviation and $|u|$ denotes the absolute value of the mean. CV is a standardized measure of dispersion, volatility, or inequality in many fields; $\mathrm{CV} = 0$ indicates zero dispersion, $x_i \equiv u$, whereas $\mathrm{CV} \gg 1$ indicates high dispersion. From (14), it follows that $\arg\min_\beta \mathbb{E}_F[\mathrm{MSE}^+]$ for a given value of CV is given by

$$\beta^*_{\mathrm{CV}} = \frac{\mathrm{CV}^2+1}{\mathrm{CV}^2+3}. \tag{18}$$

We define the efficiency of the crowd

$$\eta(\beta) \equiv 1 - \frac{\mathbb{E}_F[\mathrm{MSE}^+(\beta)]}{\mathrm{MSE}} \tag{19}$$

i.e., the relative decrease in MSE. In Fig. 4, we plot the efficiency $\eta(\beta)$ as a function of β for different values for CV. As β increases from zero, the efficiency $\eta(\beta)$ first increases, reaches its peak value, and then starts to decline. The efficiency $\eta(\beta)$ continues to increase when $\mathrm{CV} \to \infty$. When $\mathrm{CV} = 0$, there is no dispersion in x, i.e., the crowd is, in fact, not a crowd, the optimal $\beta^* = 1/3$. Both the optimal β^*_{CV} and the corresponding optimal efficiency $\eta(\beta^*_{\mathrm{CV}})$ are the increasing functions of CV or, equivalently, the dispersion. This is consistent with the wisdom of crowds literature [10]–[12], which finds that the wisdom of crowds is an increasing function of the diversity of the crowd. When $\mathrm{CV} \to \infty$, the advice $u \approx x^* = 0$; therefore, agents can safely copy the advice by setting $\beta^*_\infty = 1$. The optimal efficiency, $\eta(\beta^*_{\mathrm{CV}}) = (1+\beta^*_{\mathrm{CV}})\eta(0)$, is a factor β^*_{CV} higher than the baseline efficiency $\eta(0)$.

In practice, CV is impossible to estimate; consequently, one cannot precisely set β^*_{CV}. Nonetheless, the following guidelines still apply.

1) When no dispersion information is available, set $\beta = 1/3$.
2) When the crowd is diverse, i.e., the dispersion among individual decisions is high and CV is expected to be large, set $\beta > 1/3$.
3) When more information about λ and γ is available, one should reapply the principle of maximum entropy to compute β^*.

Our theory suggests that information about the dispersion is equally important to that about the mean. This implies the possibility of designing new behavioral research experiments or wisdom of crowds mechanisms that provide participants with both mean and dispersion.

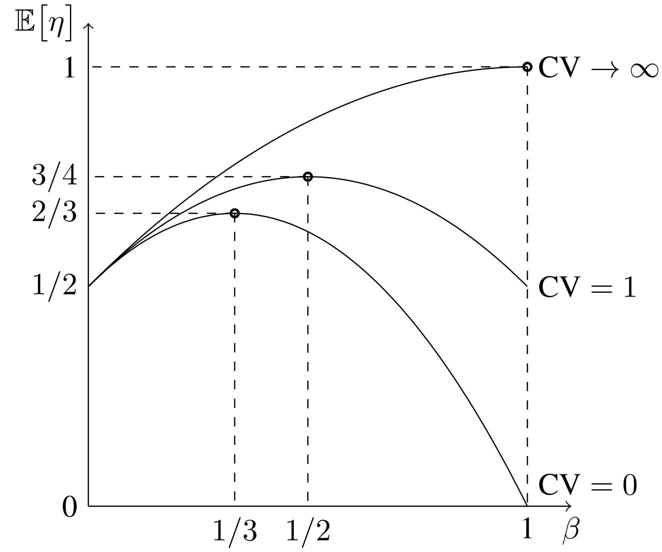

Fig. 4. $\mathbb{E}[\eta]$ as a function of β, given different CV values.

Almost any mechanism designed with a good intention also inevitably bears unintended negative consequences. Polling and other opinion-aggregating systems involving social feedback, for example, suffer from drawbacks, such as polarization [48] and data incest [49]. If we consider the game-theoretic solution to the optimization problem in (8), $\beta^* = 1/3$ might not always be the best response strategy, and free riders can set high β values to only poll opinions without contributing new information to the poll. Such a game-theoretic approach is beyond the scope of this paper to identify the minimal constructs that justify the one-third advice rule and will be investigated in the future research.

REFERENCES

[1] J. B. Soll and R. P. Larrick, "Strategies for revising judgment: How (and how well) people use others' opinions," *J. Exp. Psychol., Learn., Memory, Cognition*, vol. 35, no. 3, pp. 780–805, May 2009.
[2] S. Bonaccio and R. S. Dalal, "Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences," *Org. Behav. Hum. Decis. Processes*, vol. 101, no. 2, pp. 127–151, Nov. 2006.
[3] Y. Luo, G. Iyengar, and V. Venkatasubramanian, "Social influence makes self-interested crowds smarter: An optimal control perspective," *IEEE Trans. Comput. Social Syst.*, vol. 5, no. 1, pp. 200–209, Mar. 2018.
[4] I. Yaniv, "Receiving other people's advice: Influence and benefit," *Org. Behav. Hum. Decis. Processes*, vol. 93, no. 1, pp. 1–13, Jan. 2004.
[5] N. Harvey and I. Fischer, "Taking advice: Accepting help, improving judgment, and sharing responsibility," *Org. Behav. Hum. Decis. Processes*, vol. 70, no. 2, pp. 117–133, May 1997.
[6] J. S. Lim and M. O'Connor, "Judgemental adjustment of initial forecasts: Its effectiveness and biases," *J. Behav. Decis. Making*, vol. 8, no. 3, pp. 149–168, Sep. 1995.
[7] I. Yaniv and E. Kleinberger, "Advice taking in decision making: Egocentric discounting and reputation formation," *Org. Behav. Hum. Decis. Processes*, vol. 83, no. 2, pp. 260–281, Nov. 2000.
[8] J. B. Soll and R. Larrick, "The 80/20 rule and revision of judgment in light of another's opinion: Why do we believe in ourselves so much," in *Proc. Annu. Meeting Soc. Judgment Decis. Making*, Los Angeles, CA, USA, 1999, pp. 1–313.
[9] M. Aurelius, *Meditations*. New York, NY, USA: Modern Library, 2003.
[10] F. Galton, "Vox populi (the wisdom of crowds)," *Nature*, vol. 75, no. 7, pp. 450–451, 1907.

[11] J. Surowiecki, *The Wisdom Crowds*. New York, NY, USA: Anchor, 2005.
[12] J. Lorenz, H. Rauhut, F. Schweitzer, and D. Helbing, "How social influence can undermine the wisdom of crowd effect," *Proc. Nat. Acad. Sci. USA*, vol. 108, no. 22, pp. 9020–9025, 2011.
[13] I. Yaniv, "Weighting and trimming: Heuristics for aggregating judgments under uncertainty," *Org. Behav. Hum. Decis. Processes*, vol. 69, no. 3, pp. 237–249, 1997.
[14] A. Frachot, P. Georges, and T. Roncalli. (2001). "Loss distribution approach for operational risk." [Online]. Available: https://ssrn.com/abstract=1032523
[15] H. E. Scarf, "A min-max solution of an inventory problem," in *Proc. Stud. Math. Theory Inventory Prod.*, 1958, pp. 201–209.
[16] E. Delage and Y. Ye, "Distributionally robust optimization under moment uncertainty with application to data-driven problems," *Oper. Res.*, vol. 58, no. 3, pp. 595–612, Jun. 2010.
[17] E. T. Jaynes, "Information theory and statistical mechanics," *Phys. Rev. J. Arch.*, vol. 106, no. 4, p. 620, 1957.
[18] E. T. Jaynes, "Information theory and statistical mechanics. II," *Phys. Rev. J. Arch.*, vol. 108, no. 2, p. 171, Oct. 1957.
[19] J. P. Martino, *Technological Forecasting for Decision Making*. Amsterdam, The Netherlands: Elsevier, 1983.
[20] J. B. L. Robinson, "Delphi methodology for economic impact assessment," *J. Transp. Eng.*, vol. 117, no. 3, pp. 335–349, May 1991.
[21] A. Sourani and M. Sohail, "The delphi method: Review and use in construction management research," *Int. J. Construct. Edu. Res.*, vol. 11, no. 1, pp. 54–76, Aug. 2014.
[22] M. H. DeGroot, "Reaching a consensus," *J. Amer. Statist. Assoc.*, vol. 69, no. 345, pp. 118–121, Mar. 1974.
[23] N. E. Friedkin and E. C. Johnsen, "Social influence and opinions," *J. Math. Sociology*, vol. 15, nos. 3–4, pp. 193–206, 1990.
[24] R. Hegselmann and U. Krause, "Opinion dynamics and bounded confidence models, analysis, and simulation," *J. Artif. Societies Social Simul.*, vol. 5, no. 3, pp. 1–33, 2002.
[25] G. Weisbuch, G. Deffuant, F. Amblard, and J.-P. Nadal, "Meet, discuss, and segregate!" *Complexity*, vol. 7, no. 3, pp. 55–63, Jan./Feb. 2002.
[26] J.-H. Cho, "Dynamics of uncertain and conflicting opinions in social networks," *IEEE Trans. Computat. Social Syst.*, vol. 5, no. 2, pp. 518–531, Jun. 2018.
[27] A. Cassidy, E. Cawi, and A. Nehorai, "A model for decision making under the influence of an artificial social network," *IEEE Trans. Comput. Social Syst.*, vol. 5, no. 1, pp. 220–228, Mar. 2018.
[28] D.-S. Lee, C.-S. Chang, and Y. Liu, "Consensus and polarization of binary opinions in structurally balanced networks," *IEEE Trans. Comput. Social Syst.*, vol. 3, no. 4, pp. 141–150, Dec. 2016.
[29] J. Lorenz, "Continuous opinion dynamics under bounded confidence: A survey," *Int. J. Mod. Phys. C*, vol. 18, no. 12, pp. 1819–1838, Dec. 2007.
[30] B. A. Ogunnaike and W. H. Ray, *Process Dynamics, Modeling, and Control*, vol. 1. New York, NY, USA: Oxford Univ. Press, 1994.
[31] K. Ogata, *Modern Control Engineering*, vol. 4. Englewood Cliffs, NJ, USA: Prentice-Hall, 2002,
[32] D. E. Seborg, D. A. Mellichamp, T. F. Edgar, and F. J. Doyle, *Process Dynamics and Control*. Hoboken, NJ, USA: Wiley, 2010.
[33] J. B. Rawlings and D. Q. Mayne, *Model Predictive Control: Theory and Design*. New York, NY, USA: Nob Hill, 2009.
[34] C. S. Carver and M. F. Scheier, "Control theory: A useful conceptual framework for personality–social, clinical, and health psychology," *Psychol. Bull.*, vol. 92, no. 1, p. 111, Jul. 1982.
[35] N. Leveson, *Engineering A Safer World: Systems Thinking Applied to Safety*. Cambridge, MA, USA: MIT Press, 2011.
[36] W. M. Trochim, D. A. Cabrera, B. Milstein, R. S. Gallagher, and S. J. Leischow, "Practical challenges of systems thinking and modeling in public health," *Amer. J. Public Health*, vol. 96, no. 3, p. 538, Mar. 2006.
[37] W. T. Powers, R. Clark, and R. M. Farland, "A general feedback theory of human behavior: Part I," *Perceptual Motor Skills*, vol. 11, no. 1, pp. 71–88, Aug. 1960.
[38] W. T. Powers, R. Clark, and R. McFarland, "A general feedback theory of human behavior: Part II," *Perceptual Motor Skills*, vol. 11, no. 3, pp. 309–323, Dec. 1960.
[39] W. T. Powers, "Feedback: Beyond behaviorism: Stimulus-response laws are wholly predictable within a control-system model of behavioral organization," *Science*, vol. 179, no. 4071, pp. 351–356, Jan. 1973.
[40] W. T. Powers and W. T. Powers, *Behavior: The Control of Perception*. London, U.K.: Aldine, 1973.
[41] Y. Luo, G. Iyengar, and V. Venkatasubramanian, "Soft regulation with crowd recommendation: Coordinating self-interested agents in sociotechnical systems under imperfect information," *PLoS ONE*, vol. 11, no. 3, Mar. 2016, Art. no. e0150343.
[42] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, "Mixing beliefs among interacting agents," *Adv. Complex Syst.*, vol. 3, no. 01n04, pp. 87–98, 2000.
[43] R. Hegselmann and U. Krause, "Truth and cognitive division of labor: First steps towards a computer aided social epistemology," *J. Artif. Societies Social Simul.*, vol. 9, no. 3, p. 10, Jun. 2006.
[44] J. A. Sniezek, G. E. Schrah, and R. S. Dalal, "Improving judgement with prepaid expert advice," *J. Behav. Decis. Making*, vol. 17, no. 3, pp. 173–190, Jul. 2004.
[45] F. Gino, "Do we listen to advice just because we paid for it? The impact of advice cost on its use," *Org. Behav. Human Decis. Processes*, vol. 107, no. 2, pp. 234–245, Nov. 2005.
[46] E. A. Abbe, A. E. Khandani, and A. W. Lo, "Privacy-preserving methods for sharing financial risk exposures," *Amer. Econ. Rev.*, vol. 102, no. 3, pp. 65–70, May 2012.
[47] J. L. Jensen, "Sur les fonctions convexes et les inégalités entre les valeurs moyennes," *Acta Math.*, vol. 30, no. 1, pp. 175–193, 1906.
[48] R. Axelrod, "The dissemination of culture: A model with local convergence and global polarization," *J. Conflict Resolution*, vol. 41, no. 2, pp. 203–226, Apr. 1997.
[49] V. Krishnamurthy and W. Hoiles, "Online reputation and polling systems: Data incest, social learning, and revealed preferences," *IEEE Trans. Comput. Social Syst.*, vol. 1, no. 3, pp. 164–179, Sep. 2014.

Yu Luo received the Ph.D. degree in chemical engineering from Columbia University, New York, NY, USA, in 2017.

His doctoral training at Columbia University with Prof. V. Venkatasubramanian and Prof. G. Iyengar focused on understanding complex multiagent systems and managing systemic risk through the lens of chemical engineering—process systems engineering (PSE) in particular—and artificial intelligence. His current research with Prof. B. A. Ogunnaike and Prof. K. H. Lee at the University of Delaware, Newark, DE, USA, employs various systems techniques to design and control the upstream biopharmaceutical process of manufacturing therapeutic monoclonal antibodies. He has systematically analyzed problems at a wide range of scales and applied PSE principles to a diverse set of systems: biological systems, systems of artificial agents, social systems, financial systems, and so on. What motivates him in conducting research in PSE is the understanding of complex systems and the profound satisfaction of understanding.

Garud Iyengar received the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, USA, in 1998.

In 1998, he joined the Department of Industrial Engineering and Operations Research, Columbia University, New York, NY, USA, where he is currently the Department Chair and teaches courses in asset allocation, asset pricing, simulation, and optimization. His current research interests include convex, robust, and combinatorial optimization, queuing networks, mathematical and computational finance, and communication and information theory.

Venkat Venkatasubramanian received the Ph.D. degree in chemical engineering (with a minor in theoretical physics) from Cornell University, Ithaca, NY, USA, in 1984.

In 1985, he joined the Department of Chemical Engineering, Columbia University, New York, NY, USA, and then moved to Purdue University, West Lafayette, IN, USA, in 1988, and returned to Columbia University in 2012, where he is currently the Samuel Ruben-Peter G. Viele Professor of Engineering. His current research interests include risk analysis and management in complex engineered systems, data science methodologies for molecular products design and discovery, and design, control, and optimization through self-organization and emergence. He has authored the book *How Much Inequality is Fair? Mathematical Principles of a Moral, Optimal, and Stable Capitalist Society* (Columbia University Press, 2017).