University of Groningen

Nudging the Aggregative Behavior of Noncooperative Agents

Shakarami, Mehran; Cherukuri, Ashish; Monshizadeh, Nima

Published in:
2020 59th IEEE Conference on Decision and Control, CDC 2020

DOI:
[10.1109/CDC42340.2020.9304041](https://doi.org/10.1109/CDC42340.2020.9304041)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2020

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Shakarami, M., Cherukuri, A., & Monshizadeh, N. (2020). Nudging the Aggregative Behavior of Noncooperative Agents. In *2020 59th IEEE Conference on Decision and Control, CDC 2020* (pp. 2579-2584). [9304041] (Proceedings of the IEEE Conference on Decision and Control; Vol. 2020-December). Institute of Electrical and Electronics Engineers Inc.. <https://doi.org/10.1109/CDC42340.2020.9304041>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Nudging the Aggregative Behavior of Noncooperative Agents

Mehran Shakarami, Ashish Cherukuri, Nima Monshizadeh

Abstract—We consider the problem of steering the aggregative behavior of a set of noncooperative price-taking agents to a desired point. Different from prevalent pricing schemes where the price is available for design, we resort to suitable “nudge” mechanisms to influence the behavior of the agents. In particular, a regulator sends a price prediction signal to the agents, based on which the agents decide on their actions. This prediction is potentially different from the *actual* price, which brings the issue of reliability. We take this into account by associating trust variables to the agents, implying that the agents do not blindly follow the prediction signal. These trust variables are updated depending on the history of the discrepancy between the actual and the predicted price. We carefully examine the resulting multi-components model and analyse its convergence properties. We show that under the proposed nudge mechanisms, the regulator gains agents’ trust fully, and the aggregative behavior provably converges to a desired set point. The effectiveness of the approach is demonstrated by numerical examples.

I. INTRODUCTION

Nudge is a concept in behavioral science and economics, which is defined as any characteristic of the choice structure that predictably changes people’s behavior, without restricting any options or exceptionally affecting economic incentives¹. Consequently, nudges are not regulations, but they are easy and cheap to avoid [1]. Due to their aspects of preserving freedom of choice and being non-intrusive, nudge policies have become popular over the last few years. The most notable example is the “Behavioural Insights Team” (known as the “Nudge Unit”) that applies nudge theory in British government, and, for instance, its most recent report concerns online behavior, harm, and manipulation [2]. Another example is *informational nudging*, defined as sending manipulated, and possibly misleading, information about options to a decision maker for altering its choices [3]. Informational nudging is studied in the context of transportation systems [4] and boundedly rational decision makers [5].

Motivated by practical applications, such as charging of plug-in electrical vehicles in a coordinated way [7] or demand side management in smart grids [8], there exist various works on altering the aggregative behavior of a population of price-taking agents. A common approach in the literature is to design the price signal. If the regulator has access to all

information of the agents, as shown in [9], a linear price with respect to the actions of the agents is sufficient to achieve a desired behavior. Since this information is not available to the regulator in practice, dynamic pricing algorithms are proposed in [9]–[11]. Although treating price as a control signal facilitates the steering of aggregative behavior, the validity of such assumption can be limited since the actual price is often affected by various factors including fixed and variable production costs and daily market conditions; see e.g. [12] in the context of power systems.

Taking into account these considerations, in this manuscript, we propose nudge mechanisms through which the regulator alters aggregative behavior of price-taking agents, without directly designing the price and without fully knowing the cost/utility functions of the agents. In our setup, the regulator sends a prediction of the price to all the agents. The agents take this prediction into account when deciding their actions, but do not blindly follow it since they are aware that the prediction signal might be different than the actual price that they will incur. To model such behavior, we associate a trust variable to each agent, which increases/decreases depending on the history of the accuracy of the communicated price prediction.

Different to informational nudging in [3]–[5], we explicitly take the issue of reliability into account by incorporating trust dynamics in our design. In other words, here the agents cross-check the validity of the communicated information. Moreover, the trust dynamics couple the price prediction dynamics to the actual price, and consequently the proposed nudge mechanisms do not simplify to conventional dynamic pricing schemes.

Contributions: We present a framework which is able to capture the multi-components model resulting from nudge mechanisms in conjunction with agents’ actions and trust dynamics. Within this framework, we design two nudge mechanisms for the regulator, termed hard and soft nudge. We show that under these mechanisms, full trust of agents is gained in finite time and the aggregative behavior of the agents converges asymptotically to a desired set point. Finally, a numerical study illustrates our results. The proofs are omitted for space reasons and will appear in [13].

The rest of the paper is organized as follows. Section II provides preliminaries. The framework is introduced in Section III. Section IV presents the proposed nudge mechanisms with convergence analysis. Simulation results are included in Section V, and Section VI gathers our conclusions.

Mehran Shakarami, Ashish Cherukuri, and Nima Monshizadeh are with the Engineering and Technology Institute, University of Groningen, 9747AG, The Netherlands, m.shakarami@rug.nl, a.k.cherukuri@rug.nl, n.monshizadeh@rug.nl.

¹Nudge was originally defined as *the minimalist intervention* in a given situation such that a desired outcome is achieved [6]. However, the Nobel prize winner, Richard Thaler, presented another definition in [1], which is more popular and is used here.

II. PRELIMINARIES

This short section presents notation and basic concepts on convex analysis and projected dynamical systems.

A. Notation

We denote the set of real and nonnegative real numbers by $\mathbb{R}$ and $\mathbb{R}_{\geq 0}$, respectively. The standard Euclidean norm is denoted by $\|\cdot\|$. The symbol $\mathbf{1}_n$ denotes the vector of all ones in $\mathbb{R}^n$. Given the vectors $x_1, \dots, x_N \in \mathbb{R}^n$, we use the shorthand notation $\text{col}(x_i) = [x_1^\top, \dots, x_N^\top]^\top$. We write $A \succ 0$ to indicate that $A = A^\top \in \mathbb{R}^{n \times n}$ is positive definite. We denote the closure of a set $\mathcal{X} \subseteq \mathbb{R}^n$ by $\text{cl}(\mathcal{X})$. For a given vector $x \in \mathbb{R}^n$ and a positive definite matrix A , we define $\|x\|_A^2 := x^\top A x$. A closed ball with center $x \in \mathbb{R}^n$ and radius $r > 0$ is denoted by $\bar{B}(x, r) := \{y \in \mathbb{R}^n \mid \|x - y\| \leq r\}$. A function $F : \mathcal{X} \rightarrow \mathbb{R}^m$ is locally Lipschitz on an open set $\mathcal{X} \subset \mathbb{R}^n$ if for any point $x \in \mathcal{X}$, there exist positive scalar r and Lipschitz constant L , both dependent on x , such that $\|F(y') - F(y)\| \leq L\|y' - y\|$ for all $y', y \in \bar{B}(x, r)$. The function F is Lipschitz on $\mathcal{X}$ if there exists a positive constant L satisfying $\|F(y') - F(y)\| \leq L\|y' - y\|$ for all $y', y \in \mathcal{X}$.

B. Convex analysis

Consider a nonempty, closed, convex set $\mathcal{X} \subseteq \mathbb{R}^n$. The map $\text{proj}_{\mathcal{X}} : \mathbb{R}^n \rightarrow \mathcal{X}$ denotes the Euclidean projection on to the set $\mathcal{X}$, i.e., $\text{proj}_{\mathcal{X}}(z) := \arg \min_{y \in \mathcal{X}} \|y - z\|$. The normal cone to $\mathcal{X}$ at a given point $x \in \mathcal{X}$ is the set $\mathcal{N}_{\mathcal{X}}(x) := \{y \in \mathbb{R}^n \mid y^\top(s - x) \leq 0, \forall s \in \mathcal{X}\}$, and the tangent cone is defined as the set $\mathcal{T}_{\mathcal{X}}(x) := \text{cl}(\cup_{y \in \mathcal{X}} \cup_{\lambda > 0} \lambda(y - x))$. The projection of a vector $z \in \mathbb{R}^n$ on to $\mathcal{T}_{\mathcal{X}}(x)$ is denoted by $\Pi_{\mathcal{X}}(x, z) := \text{proj}_{\mathcal{T}_{\mathcal{X}}(x)}(z)$. Given any point $x \in \mathcal{X}$, it follows from Moreau's decomposition theorem [14, Thm. 3.2.5] that any vector $z \in \mathbb{R}^n$ can be written as $z = \text{proj}_{\mathcal{N}_{\mathcal{X}}(x)}(z) + \text{proj}_{\mathcal{T}_{\mathcal{X}}(x)}(z)$.

C. Projected dynamical systems

Given a nonempty closed set $\mathcal{X} \subset \mathbb{R}^n$ and a continuous function $h : \mathbb{R}^n \times [0, \infty) \rightarrow \mathbb{R}^n$, the nonautonomous projected dynamical system associated to them is

$$\dot{x} = \Pi_{\mathcal{X}}(x, h(x, t)). \quad (1)$$

The right-hand side of this system is in general discontinuous on the boundary of the set $\mathcal{X}$. Following [15, Def. 2.5], we specify a notion of solution to the above projected dynamical system. A map $x : [0, \infty) \rightarrow \mathcal{X}$ is a Carathéodory solution of the projected dynamical system (1) if it is absolutely continuous and satisfies $\dot{x}(t) = \Pi_{\mathcal{X}}(x(t), h(x(t), t))$ for almost all $t \in [0, \infty)$.

III. PROBLEM FORMULATION AND THE MODEL

We consider a set of agents $\mathcal{I} := \{1, \dots, N\}$ that interact repeatedly with a central regulator. Each agent is associated with a cost function J_i that it wishes to minimize by choosing

its action. In particular, the cost function of each agent $i \in \mathcal{I}$ is given by

$$J_i(x_i, p) := \frac{1}{2} (x_i - c_i)^\top Q_i (x_i - c_i) + x_i^\top p, \quad (2)$$

where $x_i \in \mathbb{R}^n$ is the action of agent i , $p \in \mathbb{R}^n$ is the price, $Q_i = Q_i^\top \in \mathbb{R}^{n \times n}$, $Q_i \succ 0$, and $c_i \in \mathbb{R}^n$. This structure appears in applications, such as charging coordination of plug-in electrical vehicles, where x_i indicates the demand of a product that comes at price p [7], [8].

Before providing further details, we give an overview of our model. The regulator provides a prediction of the price for all the agents. This prediction is potentially different from the *actual* price, which is treated as an exogenous signal and is not available for design. The agents use the price prediction to choose their action with the aim of minimizing their cost (negative of their utility) and the expected costs that they might incur under the actual price. The actual price is determined and revealed only after the actions are chosen.

The regulator, on the other hand, aims at steering *the aggregative behavior* of the agents to a desired point using the price prediction signal. We assume that the regulator does not know the cost functions of the agents. A common approach of steering aggregate behavior, often referred to as dynamic pricing, is to use the price as a control signal to regulate the system of agents [9]–[11]. In contrast, here the actual price signal is not available for design and the regulator needs to rely on the price prediction signal to manipulate the agents' behavior. Our motivation stems from the fact that, in reality, the actual price may not be prescribed a priori as a (dynamic) function of demands/actions.

Decoupling the price prediction from the actual one readily brings the issue of *trust* or *reliability*. We take this into account by considering that the agents associate a level of trust/reliability to the regulator's prediction based on the history of its accuracy.

Next, we aim to carefully model the above described features and design update schemes, termed *nudge mechanisms*, for the predicted price provided by the regulator that successfully allow steering of the aggregative behavior. We first look at the problem from agents' perspective and put forward a model where agents use available information to decide on their actions. The regulator side is considered afterwards, where *nudge mechanisms* are proposed to manipulate the aggregative behavior of the agents. We study the convergence properties of the system formed by the interconnection of agents and regulator side updates.

A. Agents' actions and trust dynamics

In choosing their actions at time $t \in [0, \infty)$, the agents have access to a price prediction $\hat{p}(t) \in \mathbb{R}^n$ sent out by the regulator. Note that this value is common for all agents. In addition, we assume that each agent $i \in \mathcal{I}$ has a local perception of the price, denoted by $\hat{\lambda}_i \in \mathbb{R}^n$, that it would have used in the absence of the prediction $\hat{p}(t)$.

As mentioned before, different from conventional dynamic pricing, the distinction between the actual price and its

prediction brings the issue of *reliability*, and we incorporate this in our model by associating a level of trust/reliability to the regulator's prediction based on the history of its accuracy. In particular, let $\gamma_i(t) \in [0, 1]$ be the trust variable of agent i associated with the price prediction $\hat{p}(t)$. Note that $\gamma_i(t) = 1$ and $\gamma_i(t) = 0$ stand for full and no trust, respectively. Given the amount of trust, predicted price, and the local perception, agent i adopts a *trust-adapted* price perception

$$\lambda_i(t) := \gamma_i(t)\hat{p}(t) + (1 - \gamma_i(t))\hat{\lambda}_i. \quad (3)$$

Note that if $\gamma_i(t)$ is close to 1, the agent disregards its own perception of the price and follows the price prediction communicated by the regulator. Conversely, as $\gamma_i(t)$ approaches 0, the agent loses trust in the price prediction $\hat{p}(t)$ and discards it when deciding on its action. The agent i uses this trust-adapted price perception to determine its action as follows:

$$x_i(t) := \arg \min_{x \in \mathbb{R}^n} J_i(x, \lambda_i(t)).$$

By using (2) and (3), the explicit expression of the action of agent i is given by

$$x_i(t) = c_i - Q_i^{-1}(\gamma_i(t)\hat{p}(t) + (1 - \gamma_i(t))\hat{\lambda}_i). \quad (4)$$

The actual price $t \mapsto p(t)$ is available to the agents once they have taken their actions. If the discrepancy between the predicted and actual price is large, then agents lose their trust in the predictions. We capture the changes of trust based on experiences by providing a trust update rule. In particular, we consider the following trust dynamics:

$$\dot{\gamma}_i(t) = \eta_i \psi_i(\|p(t) - \hat{p}(t)\|) \quad (5)$$

where $\eta_i > 0$ and $\psi_i : \mathbb{R}_{\geq 0} \rightarrow [-1, 1]$ determines whether the agent loses or gains trust in the price prediction. We assume that $\psi_i(\cdot)$ satisfies the following assumption.

Assumption III.1. The function $\psi_i : \mathbb{R}_{\geq 0} \rightarrow [-1, 1]$ is locally Lipschitz and strictly decreasing. In addition, we have $\psi_i(0) = 1$ and $\psi_i(\delta_i) = 0$ for some $\delta_i > 0$. •

The scalar δ_i quantifies the *tolerance* of agent i towards the prediction error. That is, if the error between the actual and the predicted price $\|p(t) - \hat{p}(t)\|$ is greater than δ_i , agent i begins losing trust in the prediction with the rate η_i . Conversely, trust increases as long as the error is within the tolerance δ_i . The rationale behind this dynamics is that, excluding the extreme cases of unconditional trust or distrust, trust can be gained or lost after several positive or negative experiences [16].

Note that trust variables are defined in the interval between 0 and 1. To respect this, we slightly revise (5) by adding the projection operator to it, namely:

$$\dot{\gamma}_i(t) = \Pi_{[0,1]}(\gamma_i(t), \eta_i \psi_i(\|p(t) - \hat{p}(t)\|)). \quad (6)$$

We note that the essence of the trust update rule remains the same as (5). The projection operator becomes active only if the bounds $\gamma_i = 0$ or $\gamma_i = 1$ are hit. In particular, if

$\gamma_i(t_1) = 1$ at some time $t = t_1$ and $\psi_i(\|p(t_1) - \hat{p}(t_1)\|)$ is positive (thus suggesting an increase in γ_i), the projection becomes active, and $\dot{\gamma}_i(t_1)$ is set to 0, thus prohibiting the trust variable to exceed its maximum value 1. An analogous scenario occurs for the case $\gamma_i(t_1) = 0$.

B. Desired aggregative behavior

The goal of the system regulator is to coordinate the agents such that they cumulatively behave in a desired fashion. Here, we are interested in regulating $\sum_{i \in \mathcal{I}} x_i(t)$, which we refer to as *aggregative behavior*. Such quantity often reflects total production or total demand depending on the application at hand. More precisely, the regulator aims to achieve

$$\lim_{t \rightarrow \infty} \sum_{i \in \mathcal{I}} x_i(t) = x^*, \quad (7)$$

for some desired setpoint $x^* \in \mathbb{R}^n$. To this end, we propose *nudge* mechanisms that can be implemented by the regulator. A mechanism is called a nudge if it influences the behavior of a group of individuals through providing indirect suggestions. We use this concept and propose mechanisms in which the system regulator manipulates the price prediction $\hat{p}(t)$ to achieve its goal, namely (7).

Recall that the actual price is considered here as an exogenous signal. In particular, we assume that it admits the form

$$p(t) = p_0 + \Delta p(t)$$

for all $t \in [0, \infty)$, where $p_0 \in \mathbb{R}^n$ is a constant base price, known to the regulator, and $\|\Delta p(t)\| \ll \|p_0\|$ accounts for price fluctuation. We make the following assumption that holds throughout the paper.

Assumption III.2. The actual price function $p : [0, \infty) \rightarrow \mathbb{R}^n$ is continuous, and its fluctuation satisfies $\|\Delta p(t)\| < \min_{i \in \mathcal{I}} \delta_i$ for all $t \in [0, \infty)$. •

Remark III.3. Note that without the objective (7), the best the regulator can do is to provide the agents with the true value of p_0 . In that case, the price prediction error amounts to $\|\Delta p(t)\|$. Therefore, the inequality constraint in Assumption III.2 simply means that the prediction error in such a (manipulation-free) case is within the error tolerances of all agents. •

The fact that the agents do not blindly follow $\hat{p}(t)$ implies that not any arbitrary aggregative behavior x^* is achievable. Next, we identify a set of aggregative behavior that our nudge mechanisms can successfully drive the agents to.

Consider Assumption III.2, and let $\bar{\delta} \in \mathbb{R}$ be chosen such that

$$0 < \bar{\delta} < \min_{i \in \mathcal{I}} \delta_i - \|\Delta p(t)\|, \quad \forall t \in [0, \infty). \quad (8)$$

Then, if $\hat{p}(t)$ belongs to the closed ball $\bar{B}(p_0, \bar{\delta})$, we get from the trust dynamics (6) that $\gamma_i(t)$ increases for all $i \in \mathcal{I}$. As a result, the regulator can gain trust of the agents on the price prediction by constraining $\hat{p}(t)$ to the ball $\bar{B}(p_0, \bar{\delta})$. Bearing this and the action of agents in (4) in mind, we define the

set of admissible x^* as:

$$\mathcal{X}^* := \left\{ x \in \mathbb{R}^n \mid x = \sum_{i \in \mathcal{I}} (c_i - Q_i^{-1} \hat{p}), \forall \hat{p} \in \bar{B}(p_0, \bar{\delta}) \right\}. \quad (9)$$

From (9), one can equivalently write

$$\mathcal{X}^* = \left\{ x \in \mathbb{R}^n \mid (x - x_0)^\top \left(\sum_{i \in \mathcal{I}} Q_i^{-1} \right)^{-2} (x - x_0) \leq \bar{\delta}^2 \right\},$$

where $x_0 := \sum_{i \in \mathcal{I}} (c_i - Q_i^{-1} p_0)$. Thus, the regulator can alter the aggregative behavior inside a compact set around x_0 , where the set depends on $\bar{\delta}$ and Q_i .

In other terms, $\mathcal{X}^*$ characterizes the set of aggregative behaviors that are potentially achievable while monotonically increasing the trust variables. Note from (8) that bigger the agents' error tolerances δ_i s, the larger can be $\bar{\delta}$ and thus $\mathcal{X}^*$.

For any $x^* \in \mathcal{X}^*$, there exists a unique $p^* \in \bar{B}(p_0, \bar{\delta})$ such that

$$x^* = \sum_{i \in \mathcal{I}} (c_i - Q_i^{-1} p^*), \quad (10)$$

or equivalently

$$p^* = \left(\sum_{i \in \mathcal{I}} Q_i^{-1} \right)^{-1} (-x^* + \sum_{i \in \mathcal{I}} c_i). \quad (11)$$

The vector p^* is an important quantity. If the agents fully trust the price prediction and the regulator communicates p^* as the prediction, then the aggregative behavior of the agents will be x^* . However, the regulator cannot directly compute p^* since it does not know the exact parameters defining individual cost functions. Moreover, trust can only be gained over time. To address these issues, we propose nudge mechanisms in the next section that can drive the price prediction $\hat{p}(t)$ to p^* at steady-state.

IV. REGULATOR'S NUDGE MECHANISMS

In this section, we present two nudge mechanisms to steer the aggregative behavior of the agents to the desired point.

A. Hard nudge mechanism

We propose the following update law

$$\dot{\hat{p}}(t) = \Pi_{\bar{B}(p_0, \bar{\delta})} \left(\hat{p}(t), \sum_{i \in \mathcal{I}} x_i(t) - x^* \right), \quad (12)$$

where $\bar{\delta}$ satisfies (8) and x^* is the desired aggregative behavior. The nudge mechanism in (12) updates the price predictions such that the error between the desired behavior and the current aggregative behavior diminishes. The projection constrains the predictions to the ball $\bar{B}(p_0, \bar{\delta})$ for all $t \in [0, \infty)$, thus we refer to (12) as *hard* nudge.

The overall system is obtained by interconnecting the proposed nudge mechanism (12) with agents' actions (4) and trust dynamics (6). The theorem below addresses convergence of the overall system.

Theorem IV.1. *Consider the agents' actions in (4) based on trust dynamics (6), and nudge mechanism (12) with $x^* \in \mathcal{X}^*$. Then, any solution $t \mapsto (\hat{p}(t), \text{col}(\gamma_i(t)))$, initialized*

as $(\hat{p}(0), \text{col}(\gamma_i(0))) \in \bar{B}(p_0, \bar{\delta}) \times [0, 1]^N$, converges to $(p^, \mathbb{1}_N)$ with p^* given by (11). Consequently, $\sum_{i \in \mathcal{I}} x_i(t)$ converges to x^* as desired.*

Remark IV.2. In case $x^* \notin \mathcal{X}^*$, it can be shown that the convergence of aggregative behavior is still guaranteed, but to a point which is different from x^* . In particular, the aggregative behavior converges to $x' \neq x^*$ given by

$$x' = \arg \min_{y \in \mathcal{X}^*} \frac{1}{2} \|x^* - y\|^2_{\left(\sum_{i \in \mathcal{I}} Q_i^{-1} \right)^{-1}}.$$

A detailed investigation of such a case is postponed to an extended version of this work. •

B. Soft nudge mechanism

While using the nudge mechanism in (12) is effective for driving the aggregative behavior of the agents to the desired point, convergence is guaranteed only if the price prediction is initialized in the set $\bar{B}(p_0, \bar{\delta})$. Now, we modify the nudge mechanism such that convergence is guaranteed for all $\hat{p}(0) \in \mathbb{R}^n$. The modified mechanism is given as

$$\dot{\hat{p}}(t) = \sum_{i \in \mathcal{I}} x_i(t) - x^* + \frac{1}{\varepsilon} \left(\text{proj}_{\bar{B}(p_0, \bar{\delta})} (\hat{p}(t)) - \hat{p}(t) \right), \quad (13)$$

where $\varepsilon > 0$ is a design parameter. The term $\sum_{i \in \mathcal{I}} x_i - x^*$ provides a suitable integral action as before to steer the aggregative behavior towards x^* . Note that, different from (12), this term is outside the projection operator. Therefore, the dynamics in (13) allows $\hat{p}$ to leave the set $\bar{B}(p_0, \bar{\delta})$, and we emphasize this by referring to (13) as *soft* nudge. Outside the set $\bar{B}(p_0, \bar{\delta})$, the term $\text{proj}_{\bar{B}(p_0, \bar{\delta})} (\hat{p}) - \hat{p}$ is nonzero with the penalty gain ε^{-1} , thus attracting the price prediction $\hat{p}$ to the set and preventing the loss of trust. The parameter ε is chosen sufficiently small such that trust variables monotonically increase and reach the value of 1 in finite time. Below we establish the convergence properties of the above soft nudge mechanism.

Theorem IV.3. *Consider the agents' actions in (4) based on trust dynamics (6), and nudge mechanism (13) with $x^* \in \mathcal{X}^*$. Then, there exists some $\varepsilon^* > 0$ such that for all $0 < \varepsilon \leq \varepsilon^*$, any solution $t \mapsto (\hat{p}(t), \text{col}(\gamma_i(t)))$, initialized as $(\hat{p}(0), \text{col}(\gamma_i(0))) \in \mathbb{R}^n \times [0, 1]^N$, converges to $(p^*, \mathbb{1}_N)$ with p^* given by (11). Consequently, $\sum_{i \in \mathcal{I}} x_i(t)$ converges to x^* as desired.*

Remark IV.4. While Theorem IV.3 guarantees existence of a sufficiently small ε^* , computing its value requires the knowledge of bounds on agent parameters c_i , Q_i , δ_i , and $\hat{\lambda}_i$. If such bounds are not available, then one can opt for the hard nudge mechanism in (12) at the cost of restricting the initial condition $\hat{p}(0)$ to the ball $\bar{B}(p_0, \bar{\delta})$. •

One practical issue in the presented framework could be robustness guarantees for the nudge mechanism. This requirement raises as (4) may not fully represent action selection process of the agents. As an example, the agents might be partially rational, rather than fully rational, in

choosing their optimal actions [17]. Although such robustness investigation is not carried out in this work, the soft nudge mechanism (13) prepares the ground for such a study.

V. SIMULATIONS

This section presents numerical results to validate performance of the proposed model and nudge mechanisms. We consider $N = 3$ agents with objective function (2) and

$$Q_1 = \begin{bmatrix} 3 & 2 \\ 2 & 4 \end{bmatrix}, Q_2 = \begin{bmatrix} 2 & 1 \\ 1 & 5 \end{bmatrix}, Q_3 = \begin{bmatrix} 4 & 3 \\ 3 & 4 \end{bmatrix},$$

$$c_1 = [5 \ 8]^\top, c_2 = [6 \ 3]^\top, c_3 = [7 \ 6]^\top.$$

The local perceptions of the price are chosen as

$$\hat{\lambda}_1 = [25 \ 15]^\top, \hat{\lambda}_2 = [10 \ 15]^\top, \hat{\lambda}_3 = [17 \ 23]^\top.$$

The trust variable of each agent is obtained from (6) with $\eta_1 = 0.3$, $\eta_2 = 0.2$, and $\eta_3 = 0.25$. In addition, we use the following function which satisfies Assumption III.1,

$$\psi_i(\|p - \hat{p}\|) = 2 \operatorname{sech}(\ln(a)\|p - \hat{p}\|/\delta_i) - 1,$$

with $a = (2 + \sqrt{3})$, $\delta_1 = 3$, $\delta_2 = 4$, and $\delta_3 = 3$. Note that $\psi_i(\cdot)$ is between -1 and 1 , and $\psi_i(\delta_i) = 2 \operatorname{sech}(\ln(a)) - 1 = 0$ as desired. As the agents may not trust the price prediction at the beginning, we set $\gamma_1(0) = 0.2$, $\gamma_2(0) = 0.1$, and $\gamma_3(0) = 0$.

For the price signal, we set $p_0 = [15 \ 20]^\top$ and consider $\Delta p(t)$ to be a random signal with $\|\Delta p(t)\| \leq 1$ for all $t \in [0, \infty)$. We first provide simulation results when the regulator uses the hard nudge (12) with $\bar{\delta} = 1.7$ and $\hat{p}(0) = p_0$ to steer the aggregative behavior to $x^* = [10 \ 6]^\top$.

To present the results, we denote the j -th elements of price prediction $\hat{p}$ and action x_i by $\hat{p}_{(j)}$ and $x_{i(j)}$, respectively. The price prediction and trust variables of the agents shown in Fig. 1 and Fig. 2 illustrate the resulting aggregative behavior. As expected, full trust of the agents is achieved and the aggregative behavior of the agents converges to the desired point. Fig. 3 demonstrates the relation between $\bar{B}(p_0, \bar{\delta})$ and the largest feasible set $B(p_0, \sigma)$ with $\sigma := \min_{i \in \mathcal{I}} \delta_i - \max_{t \in \mathbb{R}_{\geq 0}} \|\Delta p(t)\|$, and shows convergence of the price prediction to p^* .

We also provide numerical results of implementing the soft nudge mechanism in (13). To illustrate global convergence of the price prediction, we consider the initial condition $\hat{p}(0) = [20 \ 20]^\top$ and $\varepsilon = 0.1$. Fig. 4 depicts the price prediction and trust of the agents, and convergence of the aggregative behavior of the agents to $x^* = [10 \ 6]^\top$ is illustrated in Fig. 5. Fig. 6 shows that the price prediction converges to an invariant set in the interior of the open ball $B(p_0, \sigma)$, which is achieved by selecting small enough ε . As a result, trust of the agents on the price prediction increases, and $\hat{p}$ converges to p^* after gaining the full trust.

VI. CONCLUSIONS

We have proposed nudge mechanisms which allow a regulator to steer aggregative behavior of price-taking agents to a desired point. This is achieved without directly designing the

Fig. 1. The hard nudge mechanism's price prediction and trust variables of the agents.

Fig. 2. Aggregative behavior of the agents resulting from the hard nudge mechanism.

Fig. 3. Relation of the hard nudge mechanism's price prediction and feasible set.

Fig. 4. The soft nudge mechanism's price prediction and trust variables of the agents.

Fig. 5. Aggregative behavior of the agents resulting from the soft nudge mechanism.

Fig. 6. Relation of the soft nudge mechanism's price prediction and feasible set.

price and without complete knowledge about the cost/utility functions of the agents. We have considered explicit trust dynamics for the agents to address reliability issues that may raise due to the potential mismatch between the actual price and a price prediction signal used by the regulator to manipulate the agents' behavior. As shown both analytically and numerically, the proposed nudge mechanisms gain the trust of the agents, and subsequently drive the aggregative behavior to the setpoint.

Future works include considering different trust dynamics, more general cost functions, and investigating convergence when the desired behavior is outside the admissible set (see Remark IV.2). Robustness analysis of the proposed soft nudge mechanism is another direction for future research.

REFERENCES

- [1] R. H. Thaler and C. R. Sunstein, *Nudge: improving decisions about health, wealth, and happiness*. New Haven: Yale University Press, 2008.
- [2] E. Costa and D. Halpern, "The behavioural science of online harm and manipulation, and what to do about it: An exploratory paper to spark ideas and debate," London, pp. 1–80, 2019.
- [3] R. Guers, C. Langbort, and D. Work, "On informational nudging and control of payoff-based learning," *IFAC Proceedings Volumes*, vol. 46, no. 27, pp. 69–74, 2013.
- [4] Y. Cheng and C. Langbort, "A model of informational nudging in transportation networks," in *55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 7598–7604.
- [5] —, "On informational nudging for boundedly rational decision makers," in *57th Conference on Decision and Control (CDC)*. IEEE, 2018, pp. 4791–4796.
- [6] J. Wilk, "Mind, nature and the emerging science of change: An introduction to metamorphology," in *Metadebates on science*. Springer, 1999, pp. 71–87.
- [7] Z. Ma, D. S. Callaway, and I. A. Hiskens, "Decentralized charging control of large populations of plug-in electric vehicles," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 1, pp. 67–78, 2011.
- [8] A.-H. Mohsenian-Rad and A. Leon-Garcia, "Optimal residential load control with price prediction in real-time electricity pricing environments," *IEEE transactions on Smart Grid*, vol. 1, no. 2, pp. 120–133, 2010.
- [9] T. Alpcan and L. Pavel, "Nash equilibrium design and optimization," in *International Conference on Game Theory for Networks*. IEEE, 2009, pp. 164–170.
- [10] S. Grammatico, "Exponentially convergent decentralized charging control for large populations of plug-in electric vehicles," in *55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 5775–5780.
- [11] K. Ma, G. Hu, and C. J. Spanos, "Distributed energy consumption control via real-time pricing feedback in smart grid," *IEEE Transactions on Control Systems Technology*, vol. 22, no. 5, pp. 1907–1914, 2014.
- [12] N. P. Padhy, "Unit commitment—a bibliographical survey," *IEEE Transactions on power systems*, vol. 19, no. 2, pp. 1196–1205, 2004.
- [13] M. Shakarami, A. Cherukuri, and N. Monshizadeh, "Steering the aggregative behavior of noncooperative agents: a nudge framework," *arXiv preprint*, 2020.
- [14] J.-B. Hiriart-Urruty and C. Lemarechal, *Convex Analysis and Minimization Algorithms I: Fundamentals*. Springer Science & Business Media, 1996, vol. 305.
- [15] A. Nagurny and D. Zhang, *Projected dynamical systems and variational inequalities with applications*. Springer Science & Business Media, 2012, vol. 2.
- [16] C. M. Jonker and J. Treur, "Formal analysis of models for the dynamics of trust based on experiences," in *European Workshop on Modelling Autonomous Agents in a Multi-Agent World*. Springer, 1999, pp. 221–231.
- [17] A. Rubinstein and C.-j. Dalggaard, *Modeling bounded rationality*. MIT press, 1998.