

# Regulating Highly Automated Robot Ecologies: Insights from Three User Studies

**Wen Shen**

University of California, Irvine  
Irvine, CA, USA  
wen.shen@uci.edu

**Alanoud Al Khemeiri**

Masdar Institute  
Abu Dhabi, UAE  
aalkhemiri@gmail.com

**Abdulla Almehrzi**

Masdar Institute  
Abu Dhabi, UAE  
aa.almehrezi@gmail.com

**Wael Al Enezi**

Masdar Institute  
Abu Dhabi, UAE  
wralenezi@gmail.com

**Iyad Rahwan**

MIT Media Lab  
Cambridge, MA, USA  
irahwan@mit.edu

**Jacob W. Crandall**

Brigham Young University  
Provo, UT, USA  
crandall@cs.byu.edu

## ABSTRACT

Highly automated robot ecologies (HARE), or societies of independent autonomous robots or agents, are rapidly becoming an important part of much of the world's critical infrastructure. As with human societies, regulation, wherein a governing body designs rules and processes for the society, plays an important role in ensuring that HARE meet societal objectives. However, to date, a careful study of interactions between a regulator and HARE is lacking. In this paper, we report on three user studies which give insights into how to design systems that allow people, acting as the regulatory authority, to effectively interact with HARE. As in the study of political systems in which governments regulate human societies, our studies analyze how interactions between HARE and regulators are impacted by regulatory power and individual (robot or agent) autonomy. Our results show that regulator power, decision support, and adaptive autonomy can each diminish the social welfare of HARE, and hint at how these seemingly desirable mechanisms can be designed so that they become part of successful HARE.

## Author Keywords

Human-agent interaction; HARE; autonomy; regulation

## INTRODUCTION

Centuries of political discourse have led to diverse philosophies for how to best govern human societies. Political opinions advocate everything from highly controlled societies (authoritarianism) to loosely controlled societies (e.g., libertarianism), and just about everything in between. Political philosophies differ with respect to the extent of power and resources given to governments, as well as the rights and autonomy of individuals in society (e.g., [11, 19, 32]).

Similar discourse is necessary in the context of highly automated robot ecologies (HARE), which are collections of independent and autonomous robots, agents, or software systems that share constrained resources. For example, it is not hard to imagine future transportation systems composed almost entirely of independent driverless cars. Other aspects of modern cities, including smart grids and smart buildings, consist of networks of autonomous robotic devices that compete for and share (potentially constrained) water and electricity resources. Similarly, investor behavior in financial markets is increasingly driven by sophisticated control algorithms. As in the governing of human societies, regulators (or *regulatory authorities*) consisting of one or more people are given resources and power to influence the behavior of these HARE. The goal of these human-agent interactions is to ensure that shared resources are effectively and appropriately utilized.

Despite similarities between regulating human and robot societies, there are also glaring differences, not the least of which is that individual and collective robot behavior is often quite distinct from human behavior. Robots and other AIs can sometimes respond to stimuli instantaneously and in mass in ways that people cannot [15]. Likewise, robots are not likely to respond identically to regulations (e.g., information, incentives, and force) as people, nor might they be afforded the same rights. Thus, given the rapid rise of HARE in modern critical infrastructure, it is important that we study interactions between HARE and regulatory authorities in order to design systems that meet societal objectives.

In this paper, we study, via three user studies, how to design systems that allow people, acting as the regulatory authority, to effectively govern HARE. As in the study of political systems, our studies analyze, under an initial set of assumptions, how simple HARE are impacted by regulatory power and individual (robot) autonomy. Results show that regulator power, decision support, and adaptive robot autonomy can each diminish the social welfare of the HARE in some conditions, and suggest how these seemingly desirable mechanisms can be used so that they become part of more successful HARE.

While these user studies and the associated analysis do not (and, indeed, cannot) provide universal or general statements

about all HARE, the intended contribution of this paper is to raise awareness of the potential pitfalls and opportunities that should be considered in the design of real-world HARE.

## INTERACTING WITH HARE

Before describing the user studies, we discuss HARE and the regulator authority's role in interacting with HARE.

### HARE

A highly automated robot ecology (HARE) is a collection of independent and autonomous robots that either share resources or participate in the same activity. Both the terms *independent* and *autonomous* deserve explanations. The robots are *independent* from each other in that they are owned by different stakeholders. No one person or organization owns all robots in the collective, as individual stakeholders decide the goals and algorithms used by their robots. This independence in ownership and design (and, hence, goals and algorithms) does not imply that the robots do not impact each other. The robots may communicate with each other. Furthermore, each robot's environment is impacted by the other robots' actions.

In a HARE, the robots are *autonomous* in that, from the regulator's perspective, they make their own decisions. A regulatory authority cannot interrupt or override the robots' decision-making algorithms without the permission of their stakeholders. However, the regulator may change the robots' environment (by supplying information, providing incentives, changing physical infrastructure, etc.) to influence them.

Each robot's behavior is determined by its control algorithm. Financial incentives and other objectives often drive stakeholders to equip their robots with sophisticated and adaptive control algorithms [1, 40, 24] designed to maximize the individual stakeholders' benefits rather than societal objectives. As such, collective behavior often fails to meet societal goals. The extent to which HARE fall short of desirable societal outcomes is known as *the price of anarchy* [17, 21, 10].

### Regulation

Since the price of anarchy can be quite high [41], regulatory authorities are established to set rules and incentives that promote system-wide stability and efficiency. For example, a transportation authority assigned to regulate driverless cars can use road structure, information, and penalties and incentives (e.g., tolls or ticketing) to promote efficient and safe traffic flow. Regulators of new-age power systems can use contracts [23], information-based interventions [33], and real-time pricing [2] to influence robotic buildings to reduce peak consumption and match electricity demand to supply. In each case, the regulatory authority can potentially reduce the price of anarchy by altering the robots' environment.

Interactions between regulatory authorities and HARE bring to mind mechanism design [12], supervisory control (SC) of multiple robots [5, 43], and systems with shared autonomy [7, 13, 26]. These research areas have both key similarities and differences with regulating HARE. We discuss each in turn.

#### *Relation to Mechanism Design*

The regulatory authority engineers the environment to encourage cooperation among robots in the HARE. Cooperation is

most easily achieved by either influencing the robots to converge to a more efficient equilibrium or to alter the scenario so that it has a unique, more efficient equilibrium. This later problem, called *mechanism design* [12] (i.e., reverse game theory), has been applied to many domains including power grids [38], financial markets [20] and transportation systems [42, 34].

The goal of mechanism design is to implement strategy-proof mechanisms (e.g., payment schemes) that incentivize agents to truthfully reveal their private information [12, 27]. Unfortunately, such mechanisms do not always exist, especially in online and dynamic settings [27, 30, 29]. Furthermore, computational complexity and privacy concerns often prohibit incentive-compatible mechanisms from being implemented even when they do exist in theory [22, 28, 9]. In HARE, necessary prior knowledge required to implement such mechanisms (e.g., the robots' state and action spaces) [12, 27], is often not immediately available, and it is often not possible to obtain this information through auctions or similar revealed-preference mechanisms in a timely fashion [25, 36, 4]. Therefore, the regulator must experiment in real time to identify interventions that produce desirable societal outcomes.

#### *Relation to Supervisory Control*

Human-HARE interactions also call to mind traditional supervisory-control (SC) systems in which an operator directs multiple (semi-autonomous) robots (e.g., [5, 43]). While we anticipate that these operators face similar challenges as regulators of HARE (e.g., situation awareness [8] and operator workload), there are critical differences. For example, robots in HARE are autonomous (level 10) from the regulator's perspective, while robots in SC systems typically operate at a lower level of automation [35]. Thus, in SC, operators can directly override or alter the robots' decision-making, algorithms, or goals. Regulators of HARE cannot.

#### *Relation to Shared Autonomy*

In HARE, system dynamics are governed by the behavior of both the regulator and the robots, consistent with the idea of *shared autonomy* (also called shared control [7]). One particularly relevant application of shared control [6] that has parallels to regulating HARE is human-swarm interaction (HSI) [16, 3]. In HSI, an operator commands or influences a set of robots that have been programmed to mimic biological swarms. To do this, each robot in the swarm is equipped with simple known (to the operator) control algorithms. Despite similarities, HARE differ from traditionally defined robot swarms in that robots in HARE are programmed by separate stakeholders. Thus, the algorithms are not likely to be known to the regulator, may be highly sophisticated, and are not guaranteed to be the same among all robots. As a result, organized, cooperative group behavior can be more difficult to achieve in HARE than in traditional robot swarms, and different interactions are likely necessary.

### Parameters

Given differences between regulating HARE and other better-studied systems, we seek to understand how and when HARE can be effectively regulated. While HARE can be parameterized in many ways (including the frequency of decision-making [37, 14] and the switching processes of system

states [31]), we study two important attributes in this paper: *regulatory power* and *robot control algorithms*.

The jurisdiction and resources given to the regulatory authority to carry out its intended functions define *regulatory power*. Regulatory power determines the interventions the regulator can use. For example, local laws determine whether a transportation authority is allowed to charge tolls, how or in what manner it can change tolls, and how it can enforce payment. Furthermore, monetary resources impact which toll systems it is able to implement and maintain. Similarly, a utility company seeking to modulate the behavior of robotic buildings is limited by laws and resources that govern, among other things, the information the utility can collect and the pricing incentives it can successfully implement.

The control algorithms employed by individual robots also play an important role in HARE. Algorithms differ along many dimensions, including the data sources utilized by the algorithm, the algorithm’s depth of reasoning, and the algorithm’s adaptivity. In this paper, we consider how the ability of people to effectively regulate HARE is impacted by the ability of the robots to learn from past experiences. We refer to algorithms that do not adapt as *simple automation*, and to those that do adapt based on past experience as *adaptive automation*.

### Assumptions

Regulating HARE is a rather vast topic that we cannot fully address in a single work. For simplicity, we assume that regulators use monetary incentives to influence robot behavior and the regulatory authority consists of a single person. We also work with simulated environments (environments with simplified dynamics but which maintain many of the important characteristics of real-world HARE) to simplify data gathering. Though not without limitations, these simplifications offer a reasonable starting point to study various aspects of regulating HARE. Future work can and should relax these simplifications.

Given these assumptions, we begin to evaluate how regulatory power and algorithm adaptivity impact people’s ability to regulate HARE via a series of user studies.

### REGULATION AND ADAPTIVITY

To begin to understand how regulatory power and robot adaptivity jointly impact people’s abilities to regulate HARE, we conducted two user studies in which participants regulated simulated HARE. In the first study, participants used tolls to manage a simple transportation system composed of autonomous driverless cars. In the second study, participants regulated robotic buildings that shared a limited water supply. Both studies were 2x3 between-subjects designs in which we varied robot adaptivity and regulatory power.

#### User Study 1 – Driverless Cars

We study a HARE composed of simulated driverless cars.

##### Scenario Overview

Simulated autonomous cars used routing algorithms to navigate through a simple transportation network (Figure 1). Cars traveled at velocities determined by the number of vehicles on a road. When traffic was below a road’s capacity, cars




Figure 1. A transportation network represented by a directed graph. Circled numbers specify road capacity.

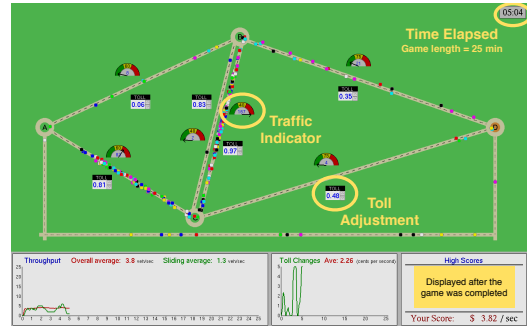


Figure 2. The GUI used in user study 1. Annotations (yellow) are overlaid for explanatory purposes.

moved at maximum speed. But when traffic approached and exceeded the road’s capacity, traffic flow slowed to a crawl (see Appendix A for details).

The regulatory authority was tasked with regulating the driverless cars so as to maximize traffic flow through the network, which was measured as the throughput through node D. To influence the cars, the regulatory authority set tolls on each road using a GUI (Figure 2) showing a bird’s-eye view of the transportation network, including the current location of each of the 300 cars. The GUI also displayed the number of cars currently on each road, as well as each road’s capacity. Toll changes were announced instantaneously to all cars.

Initially, tolls on all roads were set to \$0.50. Participants could increase or decrease each toll (between \$0.00 and \$0.99) by clicking on the corresponding buttons. Thus, if road BC was overcrowded, a regulator might consider trying to reduce the traffic congestion on this road by increasing the toll on BC, decreasing the toll on BD, increasing the toll on AB, decreasing the toll on AC, or using some combination of these methods. By properly balancing the various tolls, the regulator could eliminate congestion, which in turn produced high throughput through node D. Participants could click the buttons in rapid succession to quickly make large toll changes.

Each robot continually moved through the transportation network, repeatedly selecting a destination node and a route to that node from its current location so as to maximize its own utility. A car received positive utility each time it arrived at its destination, but incurred costs for tolls incurred and a (per unit time) operational cost. Thus, routes expected to take longer to

**Table 1. Factor levels for robot adaptivity, which were defined based on how travel costs ( $c_t(i, g)$ ) were estimated (Appendix B).**

Level	Cost Estimation
Simple automation	The cars did not learn from their past experiences. Travel costs were estimated assuming no congestion.
Adaptive automation	All cars used reinforcement learning (based on their own experiences) to determine travel costs.

traverse or that had higher tolls tended to yield lower utility and were more likely to be avoided by the cars.

Formally, each car estimated its current utility for going to destination  $g$  from its current location  $i$  as follows:

$$u(i, g) = v(g) - c_t(i, g) - c_s(i, g), \quad (1)$$

where  $v(g)$  was the utility for arriving at destination  $g$ ,  $c_t(i, g)$  was the estimated travel cost for going from the car's current location to destination  $g$ , and  $c_s(i, g)$  was the projected toll charge for going to destination  $g$  (see Appendix B for details).

Since neither  $v(g)$  nor  $c_t(i, g)$  (and how they might compare to  $c_s(i, g)$ ) were known to the regulator (for any car), the regulator could only determine how tolls might impact the cars' behavior through experimentation and observation.

#### Experimental Setup

We conducted a user study in which people regulated 300 simulated cars. In this study, we varied both *robot adaptivity* and *regulatory power* to determine how these two variables jointly impact people's ability to effectively regulate HARE. As summarized in Table 1, robot adaptivity contained two factor levels indicating the type of navigation system used by all the cars: simple automation and adaptive automation. In both cases, each car used Dijkstra's Algorithm and Eq. (1) to determine which path to follow. However, the cars used different mechanisms to estimate travel costs ( $c_t(i, g)$ ). Cars that used simple automation estimated travel costs assuming a congestion-free network. On the other hand, cars that used adaptive automation estimated travel costs on each road using reinforcement learning (Appendix B). Thus, cars that used simple automation did not learn from their past experiences (and, hence, only reacted to toll changes), whereas cars that used adaptive automation learned over time.

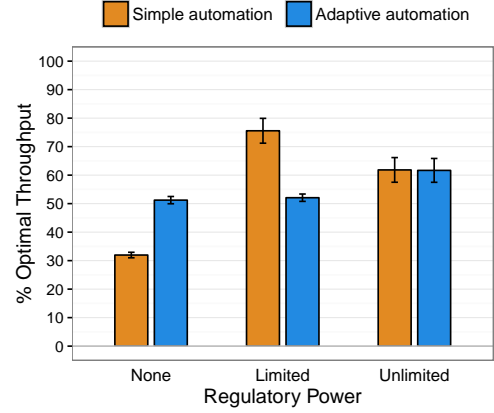
We considered three levels of regulator power: none, limited and unlimited (Table 2). For no regulatory power, no toll changes were permitted (no participants needed). When given unlimited regulatory power, participants could change tolls as frequently and as much as they desired. However, under limited regulatory power, participants were given a budget which limited the total amount of toll changes. Initially, participants received a toll-change fund of \$0.30, which increased by \$0.007 each second. Thus, the total toll-change budget for a 25-minute game was \$10.80. The absolute value of each toll change was subtracted from the budget. Toll changes were not permitted that caused the budget to drop below zero.

#### Protocol

Forty-eight students and research staff from Masdar Institute participated in the study. The following protocol was followed:

**Table 2. Factor levels for regulatory power.**

Level	Description
None	No toll changes were allowed.
Limited	Regulators had a budget which limited the amount of toll changes they could make.
Unlimited	Regulators could change tolls as much as they desired.




**Figure 3. Average throughput observed in user study 1. Error bars show the standard error of the mean.**

- The participants were randomly and uniformly assigned across four conditions: Simple-Limited, Adaptive-Limited, Simple-Unlimited, or Adaptive-Unlimited.
- The participant was trained on how to play the game in the designated condition, but with cars that chose routes randomly. This training continued until the participant felt comfortable with the objectives of the game, the user interface, and how to set tolls.
- The participant played a 25-minute game. Initially, the cars were randomly distributed across the four nodes in the network, which immediately caused congestion to develop on several roads. The participant needed to bring the system to a congestion-free state as quickly as possible. Cars were biased so that more cars preferred node C as a destination. To incentivize high performance, a high-score list was displayed once the game completed.
- The participant completed a post-experiment questionnaire, which asked which node more cars preferred and whether or not the cars employed learning algorithms.

Twelve trials for both the Simple-None and Adaptive-None conditions were also carried out (no participants required).

#### Results

Figure 3 shows the average performance of the HARE, measured as a percentage of optimal throughput over the duration of the game, achieved in each condition. Absent regulations, societies of driverless cars equipped with adaptive automation performed much better than societies of cars using simple automation. However, limited regulatory power reversed this trend. Limited regulatory power led to vastly better outcomes for societies composed of simple robots, but had no impact on societies comprised of adaptive robots. While additional (unlimited) regulatory power improved the efficiency of adaptive societies by a small amount, it decreased throughput for societies comprised of simple robots.



**Figure 4. Data from user study 1. (a) Toll adjustments made per second. (b) The number of participants that correctly deduced in the post-experiment questionnaire (top) the node most preferred by the cars and (bottom) whether or not the autonomous cars were learning.**

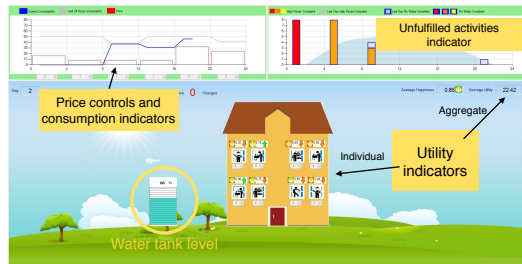
An analysis of variance, where throughput was the dependent variable and robot adaptivity and regulatory power were independent variables, confirmed many of these trends. This analysis showed a main affect for regulatory power ( $F(1,66) = 30.47, p < 0.001$ ), but not for robot adaptivity ( $F(2,66) = 0.32, p = 0.572$ ). However, there was an interaction affect between robot adaptivity and regulatory power ( $F(2,66) = 23.15, p < 0.001$ ). Tukey post hoc analysis showed that simple automation with no regulation was worse than all other conditions ( $p < 0.001$ ), while simple automation with limited regulatory power was better than all other conditions ( $p \leq 0.03$  for each pairing). Regulatory power had no significant impact on societies of adaptive robots.

We attribute the unanticipated drop in performance between the Simple-Limited to the Simple-Unlimited conditions to overuse of regulatory resources, which in turn led to participants having poorer models of the HARE. To see this, consider Figure 4a, which shows the amount of toll adjustments made by participants per second in the first user study. Unsurprisingly, substantially more toll adjustments were made by regulators who had unlimited regulatory power. While additional toll adjustments may have been justified in the case of adaptive automation, additional interventions were unnecessary when robots used simple automation. While in the Simple-Limited conditions participants were forced to wait before making more toll changes due to a limited budget, many participants did not do so in the Simple-Unlimited condition. Rather, they continually made toll adjustments without waiting sufficient time for the robots to adjust [39]. Thus, they were largely unable to effectively identify which node more robots preferred (Figure 4b-top) and whether or not the robots were learning (Figure 4b-bottom). Thus, limited resources appear to have encouraged observation and were, hence, beneficial.

In summary, moderate levels of regulatory power combined with non-adaptive robots had the highest social welfare. We now consider a second scenario to get a second data point.

### User Study 2 – Robotic Buildings

In this study, participants regulated the activity of tenants in a robotic buildings that shared a limited water supply.



**Figure 5. The (annotated) GUI used in study 2.**

### Scenario Overview

Eight (simulated) tenants of an apartment building shared a limited water resource. Each tenant’s apartment was equipped with robotic devices that automatically scheduled and executed water-related activities (e.g., laundry, dish-washing, etc.) on behalf of the tenant. A tenant programmed its own devices to execute activities automatically using a control algorithm. Water supplied to the building was collected and purified via a renewable-energy source, a process that limited water availability such that water needs exceeded supply (Appendix C).

The regulator’s job was to set the per-unit cost of water in each time period (we assumed a day with six time periods) each day such that the aggregate utility across all tenants, days, and periods was maximized. Participants set prices using the GUI pictured in Figure 5, which, in addition to allowing participants to change prices, displayed the current water level, the amount of water consumed per period, the number and value of tasks shed by the robotic devices, and the aggregate and individual happiness of the tenants.

Each tenant employed a control algorithm designed to maximize its total utility. The water needs of each tenant were defined by a set of activities. Activity  $i$  was defined by the 4-tuple  $(t_s(i), t_f(i), s(i), v(i))$ , where the time interval  $[t_s(i), t_f(i)]$  defined the time window during which activity  $i$  could be executed,  $s(i)$  was the amount of water consumed by activity  $i$ , and  $v(i)$  was how much the tenant valued the completion of activity  $i$ . When activity  $i$  was carried out, the tenant received utility  $u(i) = v(i) - c(i)$ , where  $c(i) = s(i)p(t)$  was the cost for executing activity  $i$  and  $p(t)$  was the per-unit cost of water set by the regulator for period  $t$ .

Since the tenants’ water-related activities (and how the utilities might compare to  $c(i)$ ) were unknown to the regulator, the regulator could only determine what prices to set through experimentation and observation.

### Experimental Setup

We considered societies in which (1) devices used simple (non-adaptive) algorithms and (2) devices used adaptive algorithms to schedule activities. As summarized in Table 3, simple algorithms executed any activity with positive utility when water was available. They did not adapt their behavior based on their experience. On the other hand, adaptive algorithms shifted their tenant’s activity schedules based on estimates of water availability and price in each time period (Appendix D) to maximize the tenant’s expected utility. We evaluated the same three levels of regulatory power as in Study 1 (Table 4).

**Table 3. Factor levels for robot adaptivity. See Appendix D for details.**

Level	Decision Making Process
Simple automation	The robotic building carried out activity $i$ if and only if $u_i > 0$ and there was sufficient water for the activity. The building did not shift activities based on experience.
Adaptive automation	The robotic building shifted water-related activities to maximize its tenant’s estimated utilities, which were based on estimated hourly prices and water availability. Estimates of hourly prices and water availability were based on observations made in previous days.

**Table 4. Factor levels for regulatory power.**

Level	Description
None	No price changes were allowed.
Limited	Participants were allowed to change prices no more than three times per day (by a single increment).
Unlimited	Participants were free to change prices as often and as much as they desired.


### Protocol

Forty students and research staff (mean age: 26) from Masdar Institute volunteered for the study. The participants were randomly and uniformly assigned to the same four conditions as in Study 1. Each participant was taught, via a slide presentation, how to play the game in the assigned condition. The participant then played the game in a practice scenario in which robot devices made choices randomly. Finally, the participant played a simulated 30-day game. Ten trials for both the Adaptive-None and Simple-None conditions were also conducted (no human subjects required).

### Results

Participants were asked to set prices so as to maximize the aggregate utility of all tenant’s over time. The average aggregate utility, plotted as a percentage of optimal utility, achieved in each condition is shown in Figure 6. As in Study 1, limited regulatory power produced higher social welfare in the case of simple, non-adaptive, automation. Unlimited regulatory power likewise produced lower aggregate utility than limited regulatory power when robots used simple automation. Both limited and unlimited regulatory power led to substantially lower performance when robots used adaptive algorithms.

Statistical analysis confirms these trends. A two-way analysis of variance, with aggregate utility over the last 5 days as the dependent variable and robot adaptivity and regulatory power as the independent variables, shows a main affect for both regulatory power ( $F(1,54) = 18.39, p < 0.001$ ), and robot adaptivity ( $F(1,54) = 53.53, p < 0.001$ ). There was also a significant interaction affect between robot adaptivity and regulatory power ( $F(2,66) = 22.91, p < 0.001$ ). Tukey post hoc analysis shows that, when robots used simple automation, limited regulatory power led to a significant improvement over no regulatory power ( $p = 0.037$ ) and unlimited regulatory power ( $p < 0.001$ ). Simple-Limited was also statistically better than Adaptive-Limited and Adaptive-Unlimited ( $p < 0.001$ ), and Simple-Unlimited was better than Adaptive-Unlimited ( $p = 0.023$ ). Finally, any regulation decreased the performance of societies of adaptive robots ( $p < 0.001$ ).

**Figure 6. Average aggregate utility during days 26-30 in user study 2 (water management). Error bars show the standard error of the mean.**

### USER STUDY 3 – SUPPORTING REGULATORS

The user studies described in the previous section evaluated two specific HARE. Interestingly, outcomes from both studies tell a similar story: high regulatory power combined with adaptive robots produced less efficient HARE. On the surface, these results are counter-intuitive, as both innovations seem to offer *more*. Theoretically, increased regulatory power gives the regulator more leverage over the HARE. In practice, too much regulatory power appears to have diverted the regulator away from effectively modeling the HARE. Similarly, adaptive control algorithms allow robots to, theoretically, adapt to each other, thus potentially moving the HARE towards cooperative solutions. In practice, it appears that the increased complexity of adaptive robots made it more difficult for participants to model (and, thus, influence) these HARE.

While both high regulatory power and adaptive robot control algorithms failed in the previous two studies, they may add value to the HARE under the right circumstances. One possibility is to assist the regulator in modeling the HARE. Thus, we next consider a third user study in which we gave the regulator automated support in the form of a warning system [18] that forecasted the future state of the HARE, and warned the regulator of potentially undesirable future events. We again consider the driverless-car scenario used in Study 1.

### Scenario Overview

We used a discrete-event simulation (DES) to forecast the future status of each road in the network. To do this, the system modeled the percentage of cars that chose each road at each node. These percentages, along with the number of cars currently on each road, were used to simulate the network 20 seconds in advance. The resulting simulation correctly predicted changes in future system states approximately 80% of the time. If the estimated number of cars on a road exceeded the road’s capacity at any time during the simulation, then the corresponding road was highlighted in red on the GUI. Similarly, if the estimated number of cars on a road was between 75-100%, the road was highlighted in yellow on the GUI.

### Experimental Setup and Protocol

The experimental setup and protocol was identical to our first user study, with the exception that participants were warned

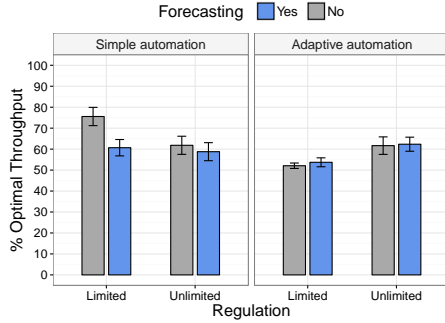


Figure 7. System throughput with and without forecasting.

Table 5. # of subjects that correctly identified the node preference and vehicle type with (Yes) and without (No) forecasting.

Condition	Node Preference		Vehicle Type	
	No / Yes	Yes / No	No / Yes	Yes / No
Simple-Limited	9 / 3	3 / 9	8 / 6	6 / 8
Simple-Unlimited	7 / 5	5 / 7	1 / 3	3 / 1
Adaptive-Limited	6 / 7	7 / 6	7 / 9	9 / 7
Adaptive-Unlimited	5 / 5	5 / 5	6 / 8	8 / 6

of pending congestion. Forty-eight participants (mean age: 28) participated this study. Twelve subjects were randomly assigned to each condition.

### Results

Figure 7 compares the average system throughput obtained when participants were given the warning system versus when they were not. In most conditions, the decision support system had little impact on the resulting performance of the HARE. The only exception was in the Simple-Limited condition, where the warning system actually appears to have *decreased* throughput. A two-way independent-samples t-test confirms this observation. In the Simple-Limited condition, the warning system significantly decreased throughput ( $M = -1.31$ ,  $SD = 1.81$ );  $t(21.743) = -2.55$ ,  $p = 0.019$ .

The post-experiment questionnaire highlights a potential explanation for the failure of the warning system. In the Simple-Limited condition, participants given the warning system had a poorer model of the HARE than those that were not given the warning system (Table 5). With the warning system, just three of the twelve participants in the Simple-Limited condition correctly identified which node more cars preferred, whereas nine of the twelve participants without the warning system correctly identified the preferred node.

Rather than learning the HARE’s tendencies, it appears that some participants instead depended on the forecasting system to identify when congestion was likely to occur. Since the warning system did not supply instructions for how to alleviate the problem [18], these participants did not know what to do once a potential problem was identified – they did not have sufficient knowledge of the HARE’s underlying tendencies.

In short, the decision-support system likewise failed to help the system take advantage of additional regulatory power and adaptive automation, and in fact appears to have made things worse. The forecasting system identified symptoms of the

underlying system, but did not help the regulator model the HARE. This negative result potentially highlights the role that a decision-support system should play in overcoming the difficulties of adaptive robots and high regulatory power. We anticipate that decision-support systems for HARE should focus on either helping the regulator to (1) form an appropriate model of robot behavior or to (2) balance the time spent modeling and implementing interventions.

### SYSTEM-SPECIFIC OR GENERAL TRENDS?

The results from the user studies reveal somewhat counter-intuitive trends about specific HARE. In particular, limited regulatory power combined with HARE with simple automation produced the best results. Would we expect these trends to generalize to other scenarios and systems, including HARE that used different adaptive control algorithms, were regulated by more or less experienced regulators, or that provided the regulator with different user interfaces? While future work is required to answer these questions in full, we seek to begin to understand the forces that impact the ability of people to regulate HARE. To do this, we use a simple mathematical model of the regulator to identify the following three general principles that appear to be influential in bringing about the results observed in our user studies.

**Principle 1:** Adaptive robot behaviors typically require the regulator to spend more time modeling the system.

**Principle 2:** Adaptive robots typically require the regulator to have higher regulatory power to effectively model the HARE.

**Principle 3:** Increased regulatory power tends to decrease the time the regulator spends modeling the HARE.

These principles appear to be applicable to all HARE, though the design of the HARE could impact the degree to which they are manifest. We discuss each in turn.

*Principle 1:* To model robot behavior, the regulator must understand how the robots will collectively react to each situation  $(s_t, h_t)$ , where  $s_t$  is the current state of the system at time  $t$  and  $h_t = (i_0, \dots, i_{t-1})$  is the intervention history the regulator has implemented up to time  $t$ . Here,  $i_k$  is the intervention carried out at time  $k$ . Let  $\mathcal{M}(s, h, i)$  describe how the robots will react when the regulator issues intervention  $i$  given system state  $s$  and intervention history  $h$ .

Since the robots’ behavior is unknown *a priori* to the regulator, the regulator must estimate  $\mathcal{M}(s, h, i)$  by observing the robots for each  $(s, h, i)$  3-tuple. The robot’s control algorithms impact the amount of time that must be given to forming the model  $\mathcal{M}(s, h, i)$ . In line with neglect benevolence [39], less time is required to model robots that use stationary decision-making processes than adaptive ones, since adaptive algorithms first adapt to the new intervention, and then react to the reactions of other robots to the intervention, and so on.

Adaptive automation also requires the operator to make more observations than simple automation. Stationary decision-making processes are typically only contingent on the current system state  $s$ , whereas adaptive ones are contingent on the tuple  $(s, h)$ . Thus, regulators must model a larger state space.

*Principle 2:* As discussed for Principle 1, adaptive algorithms require regulators to model the function  $\mathcal{M}(s, h, i)$  rather than the simpler function  $\mathcal{M}(s, i)$ . Since this model is constructed by observations that require the regulator to implement some intervention  $i$ , regulators of HARE in which robots use adaptive algorithms must have more regulatory resources to implement the necessary interventions.

*Principle 3:* More regulatory power means that regulators (a) select from a larger set of possible interventions and (b) have the ability to implement a greater number of interventions. Having more options can obviously be beneficial, but it comes at the cost of requiring the regulator to spend more time finding the best intervention among all its choices. Furthermore, implementing a greater number of interventions takes more of the regulator’s time (e.g., Figure 4a). Since the regulator must divide its time between modeling the HARE, computing effective interventions, and implementing these interventions, both of these trends mean that more regulator power can reduce the amount of time the regulator spends modeling the system. This, in turn, can lead to a poorer model of the HARE, as was observed in study 1 (Figure 4b).

The forces introduced as Principles 1-3 do not necessarily mean that adaptive control algorithms or more regulator power are always bad. We anticipate that both developments can still add value when measures are taken to counteract these forces. Future work should identify how to best do so.

## CONCLUSION

In this paper, we have presented and discussed the results of three user studies in which people regulated simulated highly automated robot ecologies (HARE). These studies provide data points that give potential insights into how we can design systems that allow people to regulate HARE so that they meet societal objectives. Though these data points only provide samples of specific HARE, they highlight easily encountered pitfalls in the design of HARE: seemingly desirable regulatory power, decision support, or adaptive robot autonomy can all lead to HARE with diminished social welfare. Our results suggest that designers of Human-HARE systems should base design decisions regarding decision support and regulatory power on helping regulators to identify and understand the underlying dynamics of the HARE rather than fixating on controlling current or future system states. Simultaneously, these data points suggest that designers of HARE should consider limiting the complexity of algorithms used by robots in the HARE, or at least to make the algorithms more immediately transparent to regulators, as simple robot autonomy coupled with limited regulatory power produced the best results.

While illuminating, we must be careful to not overstate the generality of these results, which were obtained for specific (simulated) systems. Varying any attribute of these systems (e.g., the skill and experience of the regulators; the algorithms, hardware, and information used by the robots; and the communication environment itself) could impact the results. Our studies are intended to begin to raise awareness of important issues and general principles that should be understood, weighed, and (where necessary) appropriately counteracted as

we design real-world HARE. Future work is needed to better understand, work with, and expound upon these principles.

## APPENDIX

**A. Studies 1 and 3 – Road Physics:** Congestion occurred when the number of cars on the road exceeded the road’s capacity. A car’s speed on road  $ij$  was  $V_{ij} \propto [1/(1 + e^{0.25(N_{ij}-C_{ij})})] + 0.1$ , where  $C_{ij}$  and  $N_{ij}$  were the capacity and the current number of cars on road  $ij$ , respectively. Thus, as traffic volume reached the road’s capacity, traffic flow slowed substantially.

**B. Studies 1 and 3 – Robot Behavior:** Each simulated car tried to maximize its estimated expected utility, which was based on Eq. (1). A new set of destination utilities  $v(g)$  for each node  $g \in \{A, B, C, D\}$  was generated randomly from a normal distribution each time a car reached its selected destination. Formally,  $v(g) = \mathcal{N}(b(g) + r[0, 1], 0.1 + 0.3(r[0, 1]))$ , where  $r[0, 1]$  denotes a uniform random selection from the interval  $[0, 1]$ , and  $b(g) = 0.6$  for  $g \in \{A, B, D\}$  and  $b(C) = 0.8$ . This created a preference across the HARE for node C.

The estimated travel cost  $c_s(i, g)$  (Eq. 1) was the sum of individual link costs along the shortest path to the destination. Let  $c_s(i, j)$  (defined for adjacent nodes  $i$  and  $j$ ) denote the estimated cost for traveling from  $i$  to  $j$ . Then, for cars using simple automation,  $c_s(i, j) = y \cdot x_{ij}$ , where  $y = 0.079$  was the operating cost per unit time,  $x_{ij} = L_{ij}/S_{ij}$ ,  $L_{ij}$  was the length of road  $ij$ , and  $S_{ij}$  was the car’s max speed.

Cars employing adaptive automation used reinforcement learning to estimate travel costs. Initially,  $c_s(i, j)$  was set as in simple automation. Thereafter, each time a car finished traversing road  $ij$ , it updated  $x_{ij}$  such that  $x_{ij} = \alpha x_{ij} + (1 - \alpha)z$ , where  $\alpha \in [0, 1]$  was chosen randomly for each car, and  $z$  was the observed time to traverse road  $ij$ .

**C. Study 2 – System Properties:** Each day was divided into six periods, and each tenant had one potential activity per time period. The water tank refilled at a variable rate throughout the day, such that the water-refill rate was defined by the vector  $w = (0, 40, 50, 60, 30, 0)$  (measured in *water units*). Since the consumers wished to consume 300 water units per day in aggregate, demand exceeded supply. Thus, the regulator needed to learn to set prices, via trial and error, so that water was available when the consumers had high-valued activities, which tended to be at the beginning and ending of the day.

**D. Study 2 – Robot Behavior** Formally, let  $L(d, h)$  be the amount of water available to a robotic building on day  $d$ , period  $h$ . Then, for day  $D$ , hour  $H$ , the robotic building estimates the water level to be  $L'(D, H) = 1/(D - 1) \sum_{d=1}^{D-1} L(d, H)$ . Additionally, let  $p(d, h)$  be the price of water on day  $d$ , period  $h$ . Then, the tenant estimates the price of water on day  $d + 1$ , period  $h$  to be  $p'(d + 1, h) = p(d, h)$ .

After the first day, adaptive automation shifted the tenant’s activities in day  $d$  so as to maximize expected utility. If  $L'(d, \tau) > s(\tau - t)$ , then let  $x(d, \tau, t) = \max(0, y(d, \tau, t))$ , where  $y(d, \tau, t) = v(\tau - t) - s(\tau - t)p'(d, t)$ . Otherwise,  $x(d, \tau, t) = 0$ . Then, the tenant’s schedule is shifted in day  $d$  by  $t^*(d) = \arg \max_{t \in [0, 5]} \sum_{\tau=1}^6 x(d, \tau, t)$  time periods.



## REFERENCES

1. S. M. Amin. 2000. Toward self-healing infrastructure systems. *IEEE Computer* 33, 8 (2000), 44–53.
2. S. Borenstein, M. Jaske, and A. Rosenfeld. 2002. *Dynamic Pricing, Advanced Metering and Demand Response in Electricity Markets*. Technical Report CSEM WP 105. UC Berkely: Center for the Study of Energy Markets.
3. D. S. Brown, M. A. Goodich, S.-Y. Jung, and S. Kerman. 2016. Two Invariants of Human-Swarm Interaction. *Journal of Human-Robot Interaction* 1, 5 (2016), 1–31.
4. C. P. Chambers and F. Echenique. 2016. *Revealed Preference Theory*. Cambridge University Press.
5. J. Y. C. Chen, M. J. Barnes, and M. Harper-Sciari. 2011. Supervisory Control of Multiple Social Robots: Human-Performance Issues and User-Interface Design. *IEEE Transactions on Systems, Man, and Cybernetics, Part C* 41, 4 (2011), 435–454.
6. J. W. Crandall, N. Anderson, C. Ashcraft, J. Grosh, J. Henderson, J. McClellan, A. Neupane, and M. A. Goodrich. 2017. Human-Swarm Interaction as Shared Control: Achieving Flexible Fault-Tolerant Systems. In *Proceedings of the 14th International Conference Engineering Psychology and Cognitive Ergonomics*, D. Harris (Ed.). 266–284.
7. J. W. Crandall and M. A. Goodrich. 2002. Characterizing efficiency of human robot interaction: A case study of shared-control teleoperation. In *Proceedings of the International Conference on Robots and Systems*. 1290–1295.
8. M. R. Endsley. 1988. Design and evaluation for situation awareness enhancement. In *Proceedings of the Human Factors Society's 32nd Annual Meeting*. 97–101.
9. J. Feigenbaum and S. Shenker. 2002. Distributed algorithmic mechanism design: Recent results and future directions. In *Proceedings of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*. 1–13.
10. A. G. Heldane and R. M. May. 2011. Systemic risk in banking ecosystems. *Nature* 469, 7330 (2011), 351–355.
11. L. S. Hsu. 2008. *The Political Philosophy of Confucianism*. Routledge.
12. L. Hurwicz and S. Reiter. 2006. *Designing Economic Mechanisms*. Cambridge University Press.
13. M. Johnson, J. M. Bradshaw, P. Feltovich, C. Jonker, B. van Riemsdijk, and M. Sierhuis. 2012. Autonomy and interdependence in human-agent-robot teams. *IEEE Intelligent Systems* 27, 2 (2012), 43–51.
14. N. Johnson, G. Zhao, E. Hunsader, H. Qi, N. Johnson, J. Meng, and B. Tivnan. 2013. Abrupt rise of new machine ecology beyond human response time. *Scientific reports* 3 (2013).
15. A. A. Kirilenko, A. S. Kyle, M. Samadi, and T. Tuzun. 2017. The Flash Crash: High Frequency Trading in an Electronic Market. *Journal of Finance* 73, 3 (2017), 967–998.
16. A. Kolling, K. Sycara, Nunnally S, and M. Lewis. 2013. Human Swarm Interaction: An Experimental Study of Two Types of Interaction with Foraging Swarms. *Journal of Human-Robot Interaction* 2, 2 (2013), 103–128.
17. E. Koutsoupias and C. H. Papadimitriou. 1999. Worst-case equilibria. In *Proceedings of the Symposium on Theoretical Aspects of Computer Science*. 404–413.
18. K. R. Laughery and M. S. Wogalter. 2006. Designing effective warnings. Reviews of human factors and ergonomics. *Reviews of Human Factors and Ergonomics* 2, 1 (2006), 241–271.
19. J. Locke. 1689. *Two Treatises of Government*. Awnsham Churchill.
20. J. K. MacKie-Mason and M. P. Wellman. 2006. Automated markets and trading agents. *Handbook of Computational Economics* 2 (2006), 1381–1431.
21. R. M. May and N. Arinaminpathy. 2010. Systemic risk: the dynamics of model banking systems. *Journal of the Royal Society Interface* 7, 46 (2010), 823–838.
22. F. McSherry and K. Talwar. 2007. Mechanism design via differential privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*. 94–103.
23. R. Meade and S. O'Connor. 2009. *Comparison of Long-Term Contracts and Vertical Integration in Decentralised Electricity Markets*. Technical Report EUI RSCAS; 2009/16. Robert Schuman Centre For Advanced Studies, Loyola de Palacio Programme on Energy Policy.
24. W. J. Mitchell. 2004. Beyond the ivory tower: Constructing complexity in the digital age. *Science* 303 (2004), 1472–1473.
25. R. B. Myerson. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
26. S. Nikolaidis, Y. X. Zhu, D. Hsu, and S. Srinivasa. 2017. Human-Robot Mutual Adaptation in Shared Autonomy. In *Proceedings of the International Conference on Human-Robot Interaction*.
27. B. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. 2007. *Algorithmic Game Theory*. Vol. 1. Cambridge University Press Cambridge.
28. N. Nisan and A. Ronen. 1999. Algorithmic mechanism design. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing*. ACM, 129–140.
29. D. C. Parkes and S. P. Singh. 2004. An MDP-Based Approach to Online Mechanism Design. In *Advances in Neural Information Processing Systems*. 791–798.
30. A. Pavan, I. R. Segal, and J. Toikka. 2009. Dynamic mechanism design: Incentive compatibility, profit maximization and information disclosure. *Profit Maximization and Information Disclosure* (2009).

31. T. Preis, J. J. Schneider, and H. E. Stanley. 2011. Switching processes in financial markets. *Proceedings of the National Academy of Sciences* 108, 19 (2011), 7674–7678.
32. A. Rand. 1957. *Atlas Shrugged*. New York: Random House.
33. W. P. Schultz, J. N. Nolan, R. B. Cialdini, N. J. Goldstein, and V. Griskevicius. 2007. The constructive, destructive, and reconstructive power of social norms. *Psychological Science* 18 (2007), 429–434.
34. W. Shen, C. V. Lopes, and J. W. Crandall. 2016. An Online Mechanism for Ridesharing in Autonomous Mobility-on-Demand Systems. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*.
35. T. B. Sheridan and W. L. Verplank. 1978. *Human and Computer Control of Undersea Teleoperators*. Technical Report. MIT Man-Machine Systems Laboratory.
36. P. Trigo and H. Coelho. 2011. Collective-intelligence and decision-making. In *Computational Intelligence for Engineering Systems*. Springer, 61–76.
37. A. Vespignani et. al. 2009. Predicting the behavior of techno-social systems. *Science* 325, 5939 (2009), 425.
38. P. Vytelingum, S. D. Ramchurn, T. D. Voice, A. Rogers, and N. P. Jennings. 2010. Trading agents for the smart electricity grid. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*. 897–904.
39. P. Walker, S. Nunnally, M. Lewis, A. Kolling, N. Chakraborty, and K. Sycara. 2012. Neglect Benevolence in Human Control of Swarms in the Presence of Latency. In *IEEE International Conference on Systems, Man, and Cybernetics*. 3009–3014.
40. F. Y. Wang. 2008. Toward a revolution in transportation operations: AI for complex systems. *IEEE Intelligent Systems* 23, 6 (2008), 8–13.
41. H. Youn, M. T. Gaster, and H. Jeong. 2008. Price of anarchy in transportation networks: efficiency and optimality control. *Physical Review Letters* 101(12):128701 (2008).
42. R. Zhang and M. Pavone. 2016. Control of robotic mobility-on-demand systems: a queueing-theoretical perspective. *The International Journal of Robotics Research* 35, 1-3 (2016), 186–203.
43. K. Zheng, D. F. Glas, T. Kanda, H. Ishiguro, and N. Hagita. 2014. Supervisory Control of Multiple Social Robots for Conversation and Navigation. *Transaction on Control and Mechanical Systems* 3, 2 (2014).