A Note on the Intervention Framework

Jaeok Park and Mihaela van der Schaar

August 2011

Abstract

In this document, we explain how the intervention framework proposed in our recent work is related to and different from the existing methods, concepts, and approaches in the literature. We also discuss how intervention can be combined with other approaches.

Content

1. Description of the Intervention Framework……………………………………………………1

2. Contract Theory…………………………………………………………………………...1

3. Mechanism Design……………………………………………………...………………2

4. Network Utility Maximization (NUM) ……………………………………………………...2

5. Pricing Schemes………………………………………………………………………..2

6. Repeated Games………………………………………………………………………...3

7. Stackelberg Games……………………………………………………………………...4

8. Correlated Equilibria……………………………………………………………………4

9. Bargaining Games………………………………………………………………………4

10. A Summary Table……………………………………………………………………..5

11. References……………………………………………………………………………..6

1. Description of the Intervention Framework

In our recent work [1]–[5], we proposed the intervention framework, which aims to regulate selfish behavior in multi-agent systems. In order to implement an intervention scheme, we need to augment a system with an intervention device that has a monitoring technology and an intervention capability. Using the monitoring technology, the intervention device obtains information about the actions taken by agents. The intervention capability allows the intervention device to affect the payoffs of agents, for example, by creating interference or congestion. As an example of intervention, we can consider the following scenario. Each agent's action is to request a share of an available bandwidth. If the intervention device observes that agents request "too much" bandwidth, then the device itself uses some of the available bandwidth, creating congestion (losses) for agents.

An intervention game is played in the following order.

1) The designer chooses an intervention rule, which determines an intervention action depending on the signal obtained by the intervention device.
2) The agents choose their actions independently and simultaneously, knowing the intervention rule used by the intervention device.
3) The intervention device observes a signal, which is realized following a probability distribution that depends on the agents' actions.
4) The intervention device chooses an intervention action according to the intervention rule.

We proposed a solution concept of intervention games, called an intervention equilibrium. An intervention equilibrium consists of an intervention rule and an action profile that maximize the expected payoff of the designer (or system performance) while satisfying the incentive constraints for the agents. One way to achieve an intervention equilibrium is that the designer commits to the intervention rule executed by the intervention device and recommends the action profile to the agents. Then, at an intervention equilibrium, the agents follow the recommended action profile in their self-interest.

2. Contract Theory

The literature has studied various methods to improve non-cooperative outcomes. One such method is to use contractual agreements. Contract theory is a field of economics that studies how economic actors form contractual agreements, covering the topics of incentives, information, and institutions [6]. Since intervention schemes aim to motivate agents to take appropriate actions, the intervention framework shares a theme as well as a formal framework with contract theory. In an intervention game, the intervention device obtains perfect or imperfect information about agents' actions depending on its monitoring technology. When monitoring is perfect, the intervention device can observe agents' actions without errors. When monitoring is imperfect, the intervention device cannot observe agents' actions but can only infer them from the signal it obtains.[1] Intervention with imperfect monitoring thus shares the same formal framework with the hidden action (or moral hazard) problem in contract theory. However, most works in contract theory deal with the principal-agent problem using monetary payment as the incentive device (see, for example, [8]). In contrast, the intervention framework focuses on the problem of regulating selfish behavior in resource sharing by using intervention within the system as the

[1] The terminology of perfect and imperfect monitoring follows that in the theory of repeated games [7].

incentive device. In other words, contract theory typically deals with monetary payment, which might be thought of as being *outside* the system, while the intervention framework focuses on intervention applied *inside* the system.

3. Mechanism Design

A mechanism design problem [9] considers a scenario where agents have private information and the designer desires to make a social choice (e.g., a resource allocation) depending on agents' private information. Without proper incentives, agents may not want to reveal their private information. The designer uses a mechanism in order to elicit private information from agents. In the mechanism design problem, agents have no actions to choose, and the designer can fully control the social choice. As a result, there is no issue of motivating agents to take appropriate actions in mechanism design. On the contrary, in the intervention framework, agents choose their actions while having no private information. The designer uses an intervention scheme in order to provide incentives for agents to choose appropriate actions, rather than incentives for agents to reveal private information truthfully. The intervention framework can be extended to a scenario where the designer has incomplete information about agents. In the intervention framework with incomplete information, the designer uses a procedure to induce agents to reveal their private information as well as to take appropriate actions. For example, we can consider a procedure where agents first send messages to the designer and then the designer chooses an intervention rule depending on the messages from agents.

4. Network Utility Maximization (NUM)

The Network Utility Maximization (NUM) framework, inspired by [10], aims to maximize the sum of utilities while agents maximize their individual utilities. The intervention framework differs from NUM in that it allows a more general form of the objective of the designer (maximizing the sum of the utilities of agents, as in NUM, is just one possibility). However, the main difference between NUM and intervention is that NUM assumes that agents are obedient and thus will take actions following the rule prescribed by the designer while intervention assumes that agents are selfish and thus will not obey the recommendations of the designer unless it is in their selfish interests to do so. The NUM framework aims to design a distributed algorithm that leads agents to a system-wide optimal operating point, using prices as congestion signals. In NUM, an agent's individual utility is assigned by the designer, and it is assumed that agents comply with the prescribed algorithm. On the contrary, the intervention framework aims to design an incentive scheme that induces selfish agents to achieve an incentive-constrained optimal operating point, using intervention as potential punishment for misbehavior.

5. Pricing Schemes

A well-studied class of incentive schemes uses pricing or taxation. Pricing schemes charge agents for their resource usage and have the potential to induce agents to take socially desirable actions by internalizing agents' external effects on others. The main difference between intervention and pricing is that intervention affects the resource usage of agents inside of the system whereas pricing affects the payoffs of agents through an outside instrument, money. Intervention schemes are robust in the sense that agents cannot avoid intervention as long as they use resources in the system, while pricing is not effective if agents can evade payments while still using resources. In addition, intervention schemes have informational advantages over pricing schemes. In order to achieve a desired outcome through an incentive scheme, the designer needs to know the impact of the incentive device on the payoffs of agents. Intervention affects the payoffs of agents through physical quantities (e.g., data rates and delay) associated with resource usage, and the impact of intervention on such physical quantities is relatively easy to measure. On the contrary, finding out the impact of pricing on agents requires the designer to know agents' subjective valuation of payments relative to physical quantities, which is difficult to measure.

6. Repeated Games

In the intervention framework, an intervention scheme is used to improve the (usually suboptimal) performance of Nash equilibrium in one-shot interaction. In settings where agents interact many times, the theory of repeated games [7] shows that outcomes different from (one-shot) Nash equilibrium can be obtained. The central idea of this theory is that agents can condition their current actions on past observations. Thus, in a repeated game, agents perform monitoring and provide incentives in a distributed way. In the intervention framework, on the other hand, a central entity, the intervention device, performs monitoring and provides incentives. Thus, repeated game strategies and intervention schemes can be considered as alternative methods to expand the set of equilibrium outcomes. In the repeated game approach, it is assumed that the system involves a stable long-term relationship among agents. However, many real-world systems, for instance, mobile and vehicular networks, do *not* involve such stable long-term relationships, making the theory of repeated games either not applicable or not useful. In contrast, intervention schemes can provide incentives successfully in a system with a frequently changing population. Moreover, repeated game strategies are constrained by the selfish behavior of agents. In particular, equilibrium strategies must guarantee that agents execute punishment and reward in the manner intended in their self-interest, which may require a complex structure of strategies. On the contrary, an intervention device does not have incentives of its own and thus can be programmed simply according to the design objective. We can think of the intervention device as an additional entity in the system that is specialized for monitoring and providing incentives, which reduces the burden on the agents in the system.

In a large-scale system, the burden of monitoring and providing incentives can be too heavy for a single intervention device, in which case we may need multiple intervention devices that can communicate with each other. There is a similar problem in a repeated game because the extent to which an agent can monitor and provide incentives is naturally limited to its locally interacting agents. In order to overcome this problem, agents may need to communicate their histories.[2] In the (one-shot) intervention framework, we assume that the actions of agents and the intervention action are fixed throughout the considered horizon.[3] We can introduce intervention in a repeated game so that agents and the intervention device monitor actions over time and adjust their actions based on past observations. In such a scenario, the intervention device can assist agents in punishing a deviating agent, thereby making punishment stronger.

[2] Agents can communicate their observations directly, or use a reputation scheme to share relevant information in their observations, as in [11].

[3] In [1] and [5], we consider adjustment processes for agents to reach a static intervention equilibrium. However, this is different from the repeated game approach.

7. Stackelberg Games

In the communications literature, Stackelberg games have been used to improve Nash equilibrium (see, for example, [12]). Stackelberg games divide players into two groups, a leader (sometimes multiple leaders) and followers, and the leader takes an action before the followers do. In an intervention game, the designer chooses an intervention rule before agents take their actions, and thus the designer can be considered as a leader and agents as followers. However, the choice of the leader in an intervention game is an intervention rule, which is a complete contingent plan for intervention actions to be taken given each possible signal about the actions of agents. Thus, intervention games require more overhead for the leader than Stackelberg games in that the intervention device needs to monitor the actions of agents before it chooses its action. However, intervention games are more suitable than Stackelberg games when the leader is not a resource user but a manager who regulates resource sharing by agents. In intervention games, intervention actions can be adjusted to the observed behavior of agents, and thus intervention can be applied only when punishment is needed. On the contrary, Stackelberg games lack such adaptivity.

8. Correlated Equilibria

Correlated equilibrium is a generalization of Nash equilibrium in which agents randomize in a correlated way rather than independently. In some circumstances, correlated equilibrium may improve the payoffs of all players compared to the best Nash equilibrium. Implementing a correlated equilibrium requires the presence of a mediator or a mediating device that randomizes following a correlated distribution and, on the basis of the realization of that randomization, recommends actions to each player confidentially. In the sense that there is a mediator/designer who recommends actions to players, correlated equilibrium and intervention equilibrium are similar. Furthermore, in both equilibrium concepts, agents follow the recommendations of the mediator because it is optimal for them to do so given the correlation/intervention scheme. The main difference between the two is that in correlated equilibrium there is no external punishment device while in intervention equilibrium intervention is used as punishment in case some agents deviate from the recommendations.

We note that, when a player has a dominant strategy, the dominant strategy is always the unique best response of the player no matter what belief he has about the play of other players. Therefore, in the prisoner’s dilemma game where defection is a dominant strategy for both players, using correlated equilibrium does not help because the set of correlated equilibria coincides with that of Nash equilibria. Using correlated equilibrium is helpful when some form of coordination is needed to improve performance as in a multiple access network, as analyzed in [13]. Intervention games can be extended in such a way that the designer can use a correlated distribution to determine the target action profile. After the target action profile is chosen, the designer makes confidential recommendations to agents, and intervention can be exerted when a signal that suggests a deviation occurs. This “correlated intervention equilibrium” will support more outcomes than correlated equilibrium without intervention and intervention equilibrium given the same intervention device.

9. Bargaining Games

In a bargaining game, players or a mediator chooses an operating point given a set of feasible payoffs and a disagreement point. A bargaining solution (e.g., Nash bargaining solution) proposes which operating point should be selected. We can formally define a bargaining problem based on an intervention game if we impose individual rationality constraints, and use the bargaining framework to derive a target operating point. Individual rationality constraints determine a disagreement point, while the set of operating points sustained by an intervention rule determines the set of feasible payoffs in a bargaining game. The intervention designer can be considered as the mediator in a bargaining problem who chooses the best operating point among incentive compatible and individually rational operating points according to his performance criterion.

10. A Summary Table

Approach	Description	Difference with intervention
Contract theory	A principal designs a compensation scheme for agents given informational asymmetry.	Contract theory mainly uses monetary payment as the incentive device. It deals with hidden information as well as hidden actions.
Mechanism design	Agents make reports to a mediator, and the mediator makes a social choice (e.g., a resource allocation) based on agents' reports.	The mediator, or the designer, aims to obtain the private information of agents, rather than inducing agents to take appropriate actions. The designer can control the social choice, and thus there is no hidden action problem. It is common to use taxation as the incentive device.
NUM	A manager prescribes a distributed algorithm that leads users to a socially optimal operating point.	Users are assumed to be obedient. That is, users behave not according to their innate incentives but following a rule prescribed by the manager.
Pricing schemes	Agents are charged for their resource usage.	Pricing schemes use an outside instrument, money, as the incentive device.
Repeated games	Players monitor each other and provide incentives.	The burden of monitoring and punishment is distributed to players. Players should have an incentive to execute punishment. Long-term relationship among players is needed.
Stackelberg games	Players are divided into a leader and followers. The leader takes an action first, and the followers take actions after observing the action of the leader.	The leader chooses an action, not a punishment plan.
Correlated equilibria	A mediator determines an action profile following a correlated distribution, and players receive confidential recommendations from the mediator before taking actions.	There is no punishment scheme executed by an external device.
Bargaining games	Players or a mediator chooses an operating point given a set of feasible payoffs and a disagreement point.	The intervention designer can be considered as the mediator in a bargaining problem. The set of

		feasible payoffs is determined by the set of sustainable action profiles. In order to capture the disagreement point, individual rationality constraints need to be imposed in an intervention game.

11. References

[1] J. Park and M. van der Schaar, “Stackelberg contention games in multiuser networks,” *EURASIP J. Advances Signal Process.*, vol. 2009, Article ID 305978, 15 pages, 2009.

[2] J. Park and M. van der Schaar, “Incentive provision using intervention,” in *Proc. IEEE INFOCOM* 2011.

[3] J. Park and M. van der Schaar, “Designing incentive schemes based on intervention: The case of imperfect monitoring,” in *Proc. GameNets* 2011.

[4] J. Park and M. van der Schaar, “The theory of intervention games for resource sharing in wireless communications,” *Tech. Rep.*, 2011, available: http://arxiv.org/PS_cache/arxiv/pdf/1101/1101.3052v2.pdf.

[5] Y. Xiao, J. Park, and M. van der Schaar, “Intervention in power control games with selfish users,” *Tech. Rep.*, 2011, available: http://arxiv.org/PS_cache/arxiv/pdf/1104/1104.1227v1.pdf.

[6] P. Bolton and M. Dewatripont, *Contract Theory*. Cambridge, MA: MIT Press, 2005.

[7] G. Mailath and L. Samuelson, *Repeated Games and Reputations: Long-run Relationships*. Oxford, U.K.: Oxford Univ. Press, 2006.

[8] B. Holmstrom, “Moral hazard in teams,” *Bell J. Econ.*, vol. 13, no. 2, pp. 324–340, Autumn 1982.

[9] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*, Oxford, U.K.: Oxford Univ. Press, 1995.

[10] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, “Rate control in communication networks: shadow prices, proportional fairness and stability,” *J. Operational Research Society*, vol. 49, pp. 237–252, 1998.

[11] Y. Zhang, J. Park, and M. van der Schaar, “Social norms for online communities,” *Tech. Rep.*, 2011, available: http://arxiv.org/ftp/arxiv/papers/1101/1101.0272.pdf.

[12] Y. A. Korilis, A. A. Lazar, and A. Orda, “Achieving network optima using Stackelberg routing strategies,” *IEEE/ACM Trans. Netw.*, vol. 5, no. 1, pp. 161–173, Feb. 1997.

[13] E. Altman, N. Bonneau, and M. Debbah, “Correlated equilibrium in access control for wireless communications,” in *Proc. Networking*, 2006, pp.173–183.