

The Health Impact Fund: A Mechanism Design Approach

IGH Discussion Paper #10
May 18th, 2014

Ishaan Nerurkar

Colton Jang



Incentives for Global Health

IGH is a nonprofit organization dedicated to developing market-based, systemic solutions to global health challenges.

Our main project, the Health Impact Fund, aims to increase access to medicines by creating additional incentives for innovation in the health sector.

Learn more at www.healthimpactfund.org.

The Health Impact Fund: A Mechanism Design Approach

Ishaan Nerurkar
LeapYear Innovations
ishaan@leapyearinnovations.com

Colton Jang
LeapYear Innovations
colton@leapyearinnovations.com

This paper is intended to stimulate discussion on how the Health Impact Fund can ensure that registered drugs achieve the greatest possible health impact.

Abstract: The Health Impact Fund (HIF) is a proposal designed to stimulate pharmaceutical innovation by rewarding the development of new medicines in proportion to their health impact. This discussion paper constructs a mechanism that guarantees truthful reporting from pharmaceutical companies and approximately maximizes production of QALYs, HIF's objective function. It defines the necessary components of a mechanism – the principal, the agents, the types, the outcomes, and the reactions – in the specific context of the pharmaceutical markets and mathematically proves the claims of truthful reporting and approximate optimization. The paper concludes with a comparison of the mechanism to the current HIF proposal.

The authors would like to thank Aidan Hollis and Thomas Pogge for their guiding insight and critiques.

Introduction

The primary objective of the HIF is to maximize the health impact of drugs produced by pharmaceutical companies, as measured in Quality Adjusted Life Years (QALYs). Currently, pharmaceutical companies are strongly incentivized to produce drugs that maximize return on investment, rather than drugs that increase the quality and duration of the lives of the most patients.

Our paper reframes this problem in terms of formal mechanism design. The theory of mechanism design is best understood as an extension of game theory. Where a game theorist takes the rules of a game as given and makes predictions about the beliefs and actions of strategic players, a mechanism designer actively sets the rules of the game with the goal of achieving an optimal outcome.

The key inputs of a mechanism design problem are:

1. A collective allocation problem involving many self-interested agents with private information
2. A measure of quality to evaluate any candidate solution (e.g. efficiency, profits or social welfare)
3. A description of the resources – informational or otherwise – held by the participants

Given these inputs, a mechanism prescribes the set of messages that participants can use to transmit information to each other and to the mechanism and specifies the decision that will be taken conditional on the messages that are sent. Once a mechanism is in place, participants effectively “play a game” where they send messages (e.g., a bid in an auction) as a function of their information. The goal is to find a mechanism with an equilibrium decision outcome (sometimes required to be unique) that is best according to the given measure of quality. A striking advantage of a mechanism design approach is the revelation principle, which states that for any Bayesian Nash equilibrium, there exists a corresponding mechanism with the same equilibrium outcome but in which players truthfully report type.

Mechanism design has practical applications in any collective decision problem. These currently include government auctions for drilling permits, telecom spectrum, and land; environmental protection policies regarding hunting, fishing, and carbon emissions; and multi-slot advertisement auctions in sponsored search markets.

The setting of the pharmaceutical market is apt for mechanism design – pharmaceutical companies are self-interested agents with private information, and the HIF has a clear objective of maximizing health impact. This paper outlines a mechanism that will guarantee approximate optimization of the HIF’s objective under certain conditions.

Specifications

1 Definitions

Our mechanism is a randomized algorithm in which a **principal** observes the **types** of **agents** and prescribes **outcomes** and **reactions** with the goal of maximizing the **objective function**. Let T represent the set of available types, let \mathcal{O} represent the set of outcomes, and let \mathcal{R} represent the set of reactions. Then the objective function $\Theta(\cdot)$ is defined by:

$$\Theta : T \times \mathcal{O} \times \mathcal{R} \rightarrow [0, 1]$$

1.1 The Principal

The HIF is the **principal**, or mechanism designer, and has the objective of maximizing QALYs. QALYs are maximized if

1. Pharmaceutical companies direct their research and development towards creating drugs with maximum health impact
2. Pharmaceutical companies increase access to these drugs by selling them at cost

1.2 The Agents

Agents are pharmaceutical companies. Our setting has N agents, where each agent $i \in N$ has a set of available **types** T_i which represents the agent’s valuation of participating in the HIF. This valuation is a function of the benefits and costs of registering drugs with the HIF, as compared to the current patent system.

1.2.1 Agents’ Valuation of HIF

The comparative advantages and disadvantages for a pharmaceutical company participating in the HIF have been introduced and analyzed in the Health Impact Fund book by Pogge and Hollis. [4]. Important advantages include a portion of the annual \$6 billion prize from HIF; a sharp increase in market size and medical data collection; and a reduction in marketing costs, litigation costs, duplicative research, and counterfeit drug production. In addition, pharmaceutical companies would benefit from better public perception and less government determination in funding allocation and profits. Important disadvantages include the commitment to sell registered drugs at cost for 10 years followed by open-licensing for generic production, an annual fee for health impact evaluation, and an expensive expansion into drug distribution to ensure proper distribution and usage. In addition, pharmaceutical companies cannot implement price discrimination or “evergreening” (making incremental changes to a drug and re-releasing for disproportionate increase in profits) and must provide their sales data to the HIF.

1.2.2 Agents’ Utility Function

Each agent has a utility function that is based on three factors: their type, the outcome determined by the HIF, and the agent’s reaction. An agent i has **true type** T_i which is private information and unknown to the HIF. Agent i also has **reported type** $t'_i \in T_i$ which is given to the HIF and may not necessarily be the agent’s true type. We denote a vector of reports by t , and a deviation by player i from t_i to t'_i as the vector (t'_i, t_{-i}) . The expected utilities for reporting type t'_i when all other agents report t'_{-i} are:

$$u(t_i, \mathcal{M}(t'_i, t'_{-i})) = \mathbb{E}_{o, \hat{\mathcal{R}} \sim \mathcal{M}(t'_i, t'_{-i})} [u(t_i, o, r_i(t_i, o, \hat{\mathcal{R}}_i))]$$

where $o \in \mathcal{O}$ is chosen by the mechanism, as is $\hat{\mathcal{R}}_i \subset \mathcal{R}$, the subset of reactions for each agent. Each player observes $\hat{\mathcal{R}}_i$, chooses reaction $r_i \in \hat{\mathcal{R}}_i$ and receives utility $u_i(t_i, o, r_i)$.

1.3 Outcomes and Reactions

1.3.1 Outcomes

After the HIF observes the reported types of all agents, the mechanism produces a set of outcomes. These outcomes are rules, requirements, and restrictions on pharmaceutical companies drawn from the reward mechanism proposed by Hollis and Pogge in the HIF book [4]. They include the duration of the reward period, licensing stipulations, price range, reward constraints, determination of synergistic impact, entry and exit options, and milestone payments. In order to promote consistency and stability, certain outcomes such as the reward period and licensing stipulations may remain constant over time. To give a concrete example, the mechanism will not dramatically change all outcomes and parameters every year because agents would be unwilling to participate in such an unpredictable mechanism. However, the mechanism must be flexible and adjust certain outcomes such as price range and reward constraints in order to maximize the objective function with respect to how many agents participate and with respect to the agents’ utility functions. Fixing these outcomes leaves the mechanism vulnerable to change over time.

1.3.2 Reactions

Each **outcome** corresponds to a subset of **reactions**, or actions taken by the agents. The **reaction** is the agent’s final decision on how many of its drugs it chooses to register with the HIF. Note that the reaction of an agent will always be to maximize its utility, so the choice of reaction is not a strategic decision in the game.

2 Guaranteed Results of the Mechanism

Drawing from the results of the 2012 paper by Nissim, Smorodinsky and Tennenholtz, “Approximately Optimal Mechanism Design via Differential Privacy”, we construct a mechanism \mathcal{M} that guarantees truthful reporting by pharmaceutical companies and approximately maximizes the objective function of the HIF [7]. At a high level, the mechanism \mathcal{M} is implemented as follows:

1. Select a parameter $0 \leq q \leq 1$.
2. With probability $(1 - q)$, select the outcome of the **exponential mechanism**, which is ϵ -dominant strategy truthful (1) and approximately maximizes the objective function (2).
3. With probability q , select the outcome of the **commitment mechanism**, which is strictly truthful. (3)

The **exponential mechanism** has parameter ϵ , $\mathcal{M}_E(T, \mathcal{O}, \Theta, r, \epsilon)$, and outputs an outcome $o \in \mathcal{O}$ with probability:

$$\frac{\exp\left(\frac{\epsilon \cdot \Theta(t, o, r)}{2 \cdot S(\Theta)}\right)}{\sum_{(o, r) \in \mathcal{O} \times \mathcal{R}} \exp\left(\frac{\epsilon \cdot \Theta(t, o, r)}{2 \cdot S(\Theta)}\right)}$$

The **commitment mechanism** \mathcal{M}^P selects an outcome $o \in \mathcal{O}$ uniformly at random, and sets the reactions to be: $\hat{\mathcal{R}}_i = \{r_i(t'_i, s, \mathcal{R}_i)\}$. By randomizing between these two subprocesses, the mechanism \mathcal{M} is strictly truthful if the parameter q is chosen to satisfy $\epsilon \leq \frac{q\gamma}{(1-q)|\mathcal{O}|}$ (4) and α -approximates $\Theta(\cdot)$ for values of α asymptotically upper bounded by $S(\Theta) \sqrt{\frac{|\mathcal{O}| \log |\mathcal{O}|}{\gamma N}}$ (5).

2.1 Necessary condition

In addition to appropriate selection of parameter q , the objective function must have low **sensitivity** – the maximum impact that each agent can have on the objective must be low. The lower the sensitivity, the better the approximation of the objective function. In practice, this is achieved by including many agents with relatively small investments. The sensitivity S of a function $\Theta(\cdot)$ is defined to be:

$$S(q) = \max_{i \in N, t \in T, t'_i \in T_i, o \in \mathcal{O}, r \in \hat{\mathcal{R}}} \left| \Theta(t, o, r) - \Theta((t_i, t'_i), o, r) \right|$$

2.2 Approximate maximization

A mechanism is defined as α -approximating the objective function Θ if $\forall t$,

$$\mathbb{E}_{o, \hat{\mathcal{R}} \sim \mathcal{M}}[\Theta(t, o, r(t, o, \hat{\mathcal{R}}))] \geq \max_{t, o, r} \Theta(t, o, r) - \alpha$$

The mechanism \mathcal{M} ensures that the final value of the objective function will be arbitrarily close to the maximum value. A proof of this claim is detailed in the appendix. This allows the HIF to provide a theoretical guarantee that their decisions will in expectation achieve near optimal health impact.

2.3 Truthful reporting

A mechanism is defined to be η -strictly dominant strategy truthful if $\forall i \in N, t_i \in T_i$ and $t'_{-i} \in T_{-i}$:

$$u(t_i, \mathcal{M}(t_i, t'_{-i})) \geq u(t_i, \mathcal{M}(t'_i, t'_{-i})) + \eta$$

For a fixed parameter η , the mechanism ensures that no agent can improve their utility by a false report. For small η , pharmaceutical companies will be dis-incentivized from falsely reporting their type as, as generating

a false report can be costly. In practice, a report is a normalization of the pharmaceutical companies' valuation for participating in the market opened by HIF.

There is however a possibility that pharmaceutical companies are unaware of their true type, or behave irrationally and provide a false report regardless of the disincentive to do so. In this case, if exact truthful reporting is an essential component of the HIF desiderata, we recommend a discussion of signaling mechanisms. We include the provision that the exponential mechanism is robust to such misreports, that is to say, deviation in the report of a single agent does not perturb the outcome of the mechanism greatly. This is a result of the differential privacy condition explained in the appendix.

2.4 Gap

The **gap** γ of a mechanism is a lower bound over all players and types of the worst case cost of misreporting:

$$\gamma = \min_{i, t_i \neq t'_i, t_{-i}} \max_{o \in \mathcal{O}} (u_i(t_i, o, r_i(t_i, o, \mathcal{R}_i)) - u_i(t'_i, o, r_i(t'_i, o, \mathcal{R}_i)))$$

Comparison to Current HIF Proposal

Hollis (2009) proposes four different mechanisms for obtaining maximum health impact with the minimum cost[3]:

1. Open licensing to generic firms for production. This is the least attractive to the agent as it loses control over how its product is manufactured, distributed, and marketed to wholesalers and pharmacies. It is also liable for negligently manufactured drugs produced by generic firms.
2. Price controls based on tendered bids to produce the drug by generic firms. This is slightly more attractive as the agent loses control only over how its product is manufactured, but retains control over distribution and marketing.
3. Price controls based on estimated production and distribution costs. This is most attractive to the agent as it maintains control over all aspects of manufacturing and distribution except maximum price. This is unattractive to the HIF because it is extremely difficult to estimate costs.
4. Hybrid approach (let companies choose between Option 1 and Option 2)

Each of these mechanisms is suitable only in specific circumstances. In order to consistently achieve high availability and low prices, Hollis proposes alternating between the mechanisms based on circumstance. However, there is no rigorous methodology for determining which mechanism should be applied in different circumstances. The current mechanism also cannot guarantee to pharmaceutical companies that they would strictly increase their utility by participating and participating truthfully.

The mechanism outlined in this paper gives mathematical justification for every outcome produced. Once approximations of agents' utility functions are constructed, the mechanism will choose between systems of open licensing, price controls, tendered bids, a combination of the three, or another approach entirely to ensure maximization of QALYs. The flexibility of the mechanism also addresses some of the key structural problems of the existing mechanism. Hollis (2009) stated that two outstanding problems of the current mechanism are adapting the HIF to fit the different pharmaceutical regulations of different countries and ensuring compliance with the rules and regulations established by the HIF. The flexibility of our mechanism allows the HIF to adjust outcomes from country to country with minimal loss of health impact, and the guarantee of truthful reporting ensures that pharmaceutical companies will not overstate their production costs or withhold intellectual property from generic firms.

Conclusion

The mechanism described in this proposal will ensure that pharmaceutical companies truthfully report their valuations for participation in the Health Impact Fund, and provides an approximate maximization of an objective function for the HIF, such as the QALY. The quality of the approximation will be dependent on the nature of the objective function, and in the case where the objective function has low sensitivity and the number of participants is large, the optimization will be near perfect.

Appendix

(1) For every \mathcal{O} and quality score q , the exponential mechanism is ϵ -dominant strategy truthful:

Fix an arbitrary t and t' and consider any outcome $o \in \mathcal{O}$. Let $Z(t) = \sum_{(o,r) \in \mathcal{O} \times \mathcal{R}} \exp\left(\frac{\epsilon \cdot \Theta(t,o,r)}{2 \cdot \mathcal{S}(\Theta)}\right)$

$$\begin{aligned} \frac{\Pr[M_E(t; \Theta, \mathcal{O}, r, \epsilon) = o]}{\Pr[M_E((t_{-i}, t'_i); \Theta, \mathcal{O}, r, \epsilon) = o]} &= \left(\frac{Z(t_{-i}, t'_i)}{Z(t)}\right) \left(\frac{\exp\left(\frac{\epsilon \cdot \Theta(t,o,r)}{2 \cdot \mathcal{S}(\Theta)}\right)}{\exp\left(\frac{\epsilon \cdot \Theta((t_{-i}, t'_i), o, r)}{2 \cdot \mathcal{S}(\Theta)}\right)}\right) \\ &= \left(\frac{Z(t_{-i}, t'_i)}{Z(t)}\right) \cdot \exp\left(\frac{\epsilon(\Theta(t,o,r) - \Theta((t_{-i}, t'_i), o, r))}{2 \cdot \mathcal{S}(\Theta)}\right) \\ &\leq \left(\frac{Z(t_{-i}, t'_i)}{Z(t)}\right) \cdot \left(\frac{\epsilon \cdot \mathcal{S}(\Theta)}{2 \cdot \mathcal{S}(\Theta)}\right) \leq \exp(\epsilon) \end{aligned}$$

From the fact that $\exp(\epsilon) \geq 1 - \epsilon$ and that the utilities are normalized to $[0, 1]$ that $\mathbb{E}_{o \sim M_E(t)}[u_i(o)] \geq \exp(-\epsilon) \mathbb{E}_{o \sim M(t_{-i}, t'_i)}[u_i(o)] \geq \mathbb{E}_{o \sim M(t_{-i}, t_i)}[u_i(o)] - \epsilon$.

(2) The exponential mechanism approximately maximizes Θ . Let OPT denote the optimal value of Θ , and $o = \mathcal{M}_E(t; \Theta, \mathcal{O}, r, \epsilon)$.

For any t ,

$$\Pr[\Theta(t, o, r, \epsilon) \leq t] \leq \frac{\Pr[\Theta(t, o, r) \leq t]}{\Pr[\Theta(t, o, r) = \text{OPT}]} \leq |\mathcal{O}| \cdot \exp\left(\frac{\epsilon(t - \text{OPT})}{2 \cdot \mathcal{S}(\Theta)}\right)$$

Let $t = \text{OPT} - \frac{2\mathcal{S}(\Theta)}{\epsilon}(\log |\mathcal{O}| + t)$. Then,

$$\Pr[\Theta(t, o) \leq t] \leq |\mathcal{O}| \exp(-(\log |\mathcal{O}| + t)) = \exp(-t)$$

Thus we have shown that with high probability, the value of Θ returned by the mechanism is within $\frac{2\mathcal{S}(\Theta)}{\epsilon}(\log |\mathcal{O}| + t)$ of the optimal.

(3) The commitment mechanism is strictly truthful For any i , let $o^* = \arg \min_{t'_i \neq t_i, t_{-i}} \max_{o \in \mathcal{O}} (u_i(t_i, o, r_i) - u_i(t_i, o, r_i))$. For any deviation t'_i ,

$$u(t_i, \mathcal{M}^P(t_i, t_{-i})) = \sum_{o \in \mathcal{O}} \frac{1}{|\mathcal{O}|} \cdot u(t_i, o, r_i) \geq \frac{1}{|\mathcal{O}|} \left(\sum_{o \neq o^*} u(t_i, o, r_i) + u(t'_i, o^*, r_i) + \gamma \right) = u(t_i, \mathcal{M}^P(t'_i, t_{-i})) + \frac{\gamma}{|\mathcal{O}|}$$

Therefore for every agent has an incentive of $\frac{\gamma}{|\mathcal{O}|}$ to report truthfully.

(4) \mathcal{M} is strictly truthful for $\epsilon \leq \frac{q\gamma}{(1-q)|\mathcal{O}|}$. For all t_i, t'_i, t_{-i} ,

$$u_i(t_i, \mathcal{M}(t_i, t_{-i})) = (1-q)u_i(t_i, \mathcal{M}_E(t_i, t_{-i})) + qu_i(t_i, \mathcal{M}^P(t_i, t_{-i})) \geq u_i(t_i, \mathcal{M}(t'_i, t_{-i})) - (1-q)\epsilon + q\frac{\gamma}{|\mathcal{O}|}$$

(5) \mathcal{M} approximately maximizes Θ . Let $q = \frac{|\mathcal{O}|\epsilon}{\gamma + |\mathcal{O}|\epsilon}$ to satisfy the truthfulness condition.

$$\mathbb{E}_{o, r \sim \mathcal{M}}[\Theta(t, o, r)] \geq (1-q)\mathbb{E}_{o, r \sim \mathcal{M}_E}[\Theta(t, o, r)] = \left(1 - \frac{|\mathcal{O}|\epsilon}{\gamma + |\mathcal{O}|\epsilon}\right) \left(\text{OPT} - O(\mathcal{S}(\Theta) \frac{\log |\mathcal{S}|}{\epsilon})\right) \geq \text{OPT} - \frac{|\mathcal{O}|\epsilon}{\gamma + |\mathcal{O}|\epsilon} - O(\mathcal{S}(\Theta) \frac{\log |\mathcal{O}|}{\epsilon})$$

References

- [1] Balcan M-F, Avrim Blum, Jason D Hartline, and Yishay Mansour. (2005) Mechanism design via machine learning. In Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on, pages 605-614. IEEE, 2005. Accessed March 18th 2014 at <http://eprints.pascal-network.org/archive/00001112/01/05-focs-aa.pdf>
- [2] Dwork, Cynthia, Frank McSherry, Kobbi Nissim, and Adam Smith. (2006) Calibrating noise to sensitivity in private data analysis. In TCC '06, pages 265-284, 2006. Accessed March 18th 2014 at <http://www.cse.psu.edu/~asmith/privacy598/papers/dmns06.pdf>
- [3] Hollis, Aidan (2009). The Health Impact Fund and Price Determination Incentives for Global Health. Accessed March 20th 2014 at http://healthimpactfund.org/wp-content/uploads/2012/11/DP1_Hollis.pdf
- [4] Hollis, Aidan and Thomas Pogge (2009). The Health Impact Fund: Making New Medicines Accessible for All. Incentives for Global Health. Forthcoming.
- [5] Legros, Patrick and Estelle Cantillon. "The Nobel Prize: What Is Mechanism Design and Why Does It Matter for Policy-making?" VoxEU.org. Vox, 18 Oct. 2007. Web. Accessed March 21st, 2014 at <http://voxeu.org/article/nobel-prize-what-mechanism-design-and-why-does-it-matter>
- [6] McSherry, Frank and Kunal Talwar. (2007) Mechanism design via differential privacy. In FOCS, pages 94-103, 2007. Accessed March 18th 2014 at <http://www.cse.psu.edu/~asmith/privacy598/papers/mt08.pdf>
- [7] Nissim, Kobbi, Rann Smorodinsky, and Moshe Tennenholtz. (2012) Approximately optimal mechanism design via differential privacy. In ITCS, pages 203-213, 2012. Accessed March 18th 2014 at <http://arxiv.org/pdf/1004.2888.pdf>