E C O N O M I C S   B U L L E T I N

Learning in Bayesian regulation: desirable or undesirable?

Semih Koray
Bilkent University

Ismail Saglam
Bogazici University

Abstract

We examine the social desirability of learning about the regulated agent in a generalized principal-agent model with incomplete information. An interesting result we obtain is that there are situations in which the agent prefers a Bayesian regulator to have more, yet incomplete, information about his private type.

The authors thank Leonid Hurwicz, Murat Sertel, and seminar participants at Bilkent University and Bogazici University. Saglam acknowledges the support of Turkish Academy of Sciences, in the framework of `Distinguished Young Scientist Award Program' (TUBA-GEBIP-2004). The final revision of this paper was made in 2007 while Saglam was visiting the Economics Department of Massachusetts Institute of Technology to which he is grateful for its hospitality. The usual disclaimer applies.
Citation: Koray, Semih and Ismail Saglam, (2007) "Learning in Bayesian regulation: desirable or undesirable?." *Economics Bulletin*, Vol. 3, No. 12 pp. 1-10
Submitted: March 6, 2007. **Accepted:** April 8, 2007.
URL: http://economicsbulletin.vanderbilt.edu/2007/volume3/EB-07C70006A.pdf

1. Introduction

The issue of learning has occupied an important place in the recent literature of game theory while most of the pioneering studies have focused on learning in repeated games with incomplete information. For example, Jordan (1991) considers a noncooperative normal form game where each player is endowed with full Bayesian rationality and has prior beliefs about his opponents' privately known payoffs. The Bayesian Nash equilibrium of this game needs not coincide with the Nash equilibrium of the complete information (true) game. However, Jordan shows that under certain restrictions on beliefs the players in a repeated play of the described normal form game can learn to play the Nash equilibrium of the complete information game even though they will not necessarily attain complete information. Kalai and Lehrer (1993) and Blume and Easley (1995) obtain a similar convergence result for infinitely repeated games that involve non-myopic players. The empirical evaluations of the Jordan's Bayesian learning model was later evaluated in Cox, Shachat and Walker (2001), which shows that when the true game had a unique pure strategy equilibrium, the experimental subjects' play converged to the equilibrium, while this was not the case if the true game had multiple equilibria.

In the existing literature, learning occurs while each player maximizes his infinite horizon expected utility and updates his prior beliefs using the Bayes rule. However, in this paper we examine the issue of Bayesian learning as a direct goal of (one of the) players in a static decision problem and ask the following questions: in a principal-agent model of regulation with incomplete information that borrows from Guesnerie and Laffont (1984), (i) what is 'more information' in a situation of 'incomplete' learning where the belief of the regulator about the regulated agent does not coincide with the truth? (ii) is 'more information' about the regulated agent always desirable for the regulator and the principal or, conversely, undesirable for the regulated agent?

The organization of the paper is as follows: Section 2 introduces the Bayesian regulation model. We present our results in Section 3. Finally, Section 4 concludes.

2. Model

Consider two players with quasi-linear utility functions

$$u_p(x, t, \theta) = V_p(x, \theta) - t, \tag{1}$$

$$u_a(x, t, \theta) = V_a(x, \theta) + t, \tag{2}$$

where V_p and V_a (u_p and u_a) stand for the utilities (net utilities) of the principal and the agent, respectively. Here, θ is the agent's private information about his utility function, x is called a decision and t is the total monetary transfer from the principal to the agent.[1]

[1] For example, in a setting of monopoly regulation, θ can be considered as the private cost parameter of a monopolist, x the output decision, V_p the consumers' surplus, and V_a the monopolist's operating profits.

The private type parameter θ of the agent is commonly known to lie in some closed interval Θ of reals. Define $\theta_0 = \min(\Theta)$ and $\theta_1 = \max(\Theta)$. We also assume that:

*A*0. $\text{argmax}_x V_p(x,\theta) \neq \text{argmax}_x V_a(x,\theta)$
*A*1. $\partial(V_p + V_a)/\partial x > 0$
*A*2. $\partial^2(V_p + V_a)/\partial x^2 < 0$
*A*3. $\partial^2 V_p/\partial x \partial\theta \leq 0$
*A*4. $\partial^2 V_a/\partial x \partial\theta \leq 0$
*A*5. $\partial V_a/\partial\theta < 0$
*A*6. $\partial^3 V_a/\partial x \partial\theta^2 \leq 0$
*A*7. $\partial^3 V_a/\partial x^2 \partial\theta \leq 0$

The regulator announces a contract between the principal and the agent. The instruments of the contract are the control of the decision x and the transfer t to the agent. By the Revelation Principle (Gibbard, 1973; Myerson, 1979), the regulator can restrict himself to direct revelation mechanisms which ask the agent to report his private information and which give to the agent no incentive to lie. The optimal regulatory policy is designed to satisfy two conditions. First, the agent must never expect a greater net utility by misreporting than he could by truthfully reporting his private information:

$$\text{(IC)} \quad u_a(x(\theta), t(\theta), \theta) \geq u_a(x(\hat{\theta}), t(\hat{\theta}), \theta), \text{ for all } \theta, \hat{\theta} \in \Theta \tag{3}$$

The second condition is that the regulator must never regulate the agent without guaranteeing him a nonnegative net utility:

$$\text{(IR)} \quad u_a(x(\theta), t(\theta), \theta) \geq 0, \text{ for all } \theta \in \Theta \tag{4}$$

Now, let $U_a(\theta, \hat{\theta})$ denote the net utility of the agent when he reports his private parameter as $\hat{\theta}$ while θ is the actual parameter. Condition (IC) implies that $U_a(\theta, \theta) = U_a(\theta)$ satisfies

$$U_a(\theta) = \max_{\hat{\theta} \in \Theta} \ u_a(x(\hat{\theta}), t(\hat{\theta}), \theta) = u_a(x(\theta), t(\theta), \theta) \tag{5}$$

for all $\theta \in \Theta$. From the envelope theorem, we obtain

$$\frac{dU_a}{d\theta} = \frac{\partial u_a}{\partial\theta} = \frac{\partial V_a}{\partial\theta}. \tag{6}$$

Similarly, denote by $U_p(\theta)$ the net utility of the principal when the agent truthfully reports his private parameter as θ.

The social welfare $W(\theta)$ is defined as the sum of the principal's net utility and a fraction of the agent's net utility:

$$W(\theta) = U_p(\theta) + \alpha U_a(\theta), \tag{7}$$

where $\alpha \in [0,1]$ is the relative weight assigned to the net utility of the agent. Integrating (6), using the assumption (A5), yields

$$U_a(\theta) = -\int_{\theta}^{\theta_1} \frac{\partial}{\partial \tilde{\theta}} V_a(x(\tilde{\theta}), \tilde{\theta}) d\tilde{\theta}. \tag{8}$$

Inserting $U_p(\theta) = V_p(x(\theta), \theta) - t(\theta)$ and $t(\theta) = U_a(\theta) - V_a(\theta)$ into (7), the actual social welfare becomes:

$$W(\theta) = V_p(x(\theta), \theta) + V_a(x(\theta), \theta) + (1-\alpha) \int_{\theta}^{\theta_1} \frac{\partial}{\partial \tilde{\theta}} V_a(x(\tilde{\theta}), \tilde{\theta}) d\tilde{\theta} \tag{9}$$

Assumptions 6 and 7 are sufficent for the optimal decision $x(.)$, if exists, to be nonincreasing and implemented by the described subsidy mechanism. However, it is known that there exists no feasible solution $x(.)$ that maximizes (9) unless the two players' welfares are equally weighted in the social welfare function or that the utility of the agent is seperable in its two arguments. The common remedy is to introduce a Bayesian regulator.

We consider a Borel field $\mathcal{T}^{\Theta}$ on the type space Θ and regard the subset $\mathcal{A}^{\Theta}$ of probability measures on $\mathcal{T}^{\Theta}$ with densities that are strictly positive at each element of Θ as the set of admissible prior beliefs for the regulator. Let $f \in \mathcal{A}^{\Theta}$ be the prior belief of the regulator and F be the respective cumulative distribution function. We assume that f becomes common knowledge before the regulator asks the agent to report his type. Let the pair (f, Θ) denote the information structure that is commonly known by all parties in the society.

The objective function of the regulator under the structure (f, Θ) is the expected social welfare:

$$\int_{\theta_0}^{\theta_1} \left(V_p(x(\theta), \theta) + V_a(x(\theta), \theta) + (1-\alpha) \int_{\theta}^{\theta_1} \frac{\partial}{\partial \tilde{\theta}} V_a(x(\tilde{\theta}), \tilde{\theta}) d\tilde{\theta} \right) f(\theta) d\theta \tag{10}$$

Modifying (10), we obtain the problem of the Bayesian regulator as:

$$\max_{x(.)} \int_{\theta_0}^{\theta_1} \left(V_p(x(\theta), \theta) + V_a(x(\theta), \theta) + (1-\alpha) \frac{F(\theta)}{f(\theta)} \frac{\partial}{\partial \theta} V_a(x(\theta), \theta) \right) f(\theta) d\theta \tag{11}$$

s.t. (IC) and (IR)

To simplify the solution and its analysis, we will assume that for all $\Theta \subset I\!R$ and $f \in \mathcal{A}^{\Theta}$:

A8. $F(\theta)/f(\theta)$ is nondecreasing in θ

Proposition 1. *The solution to Bayesian regulation problem (11) satisfies*

$$\frac{\partial V_p}{\partial x} + \frac{\partial V_a}{\partial x} = -(1-\alpha) \frac{F(\theta)}{f(\theta)} \frac{\partial^2 V_a}{\partial x \partial \theta}. \tag{12}$$

We henceforth assume $\alpha \in [0,1)$ and $\partial^2 V_a / \partial x \partial \theta < 0$ in order to be in the Bayesian framework where the beliefs of the regulator affects the optimal program (12) through the term $F(\theta)/f(\theta)$, so called "the inverse of the reverse hazard rate".

Let $\bar{x}^f$ denote the solution to (12), and let $\bar{U}_p^f(\theta), \bar{V}_p^f(\bar{x}^f(\theta), \theta), \bar{U}_a^f(\theta)$, $\bar{V}_a^f(\bar{x}^f(\theta), \theta)$, $\bar{t}^f(\theta)$, and $\bar{W}^f(\theta)$ respectively denote the net and gross utilities of the principal and the agent, the subsidy and the social welfare at the report $\theta \in \Theta$ under the belief $f(.)$.

3. Results

We first define a dominance relation over the set of admissible beliefs to compare the regulatory outcomes that these beliefs lead to.

Definition 1. Let $f_1 \in \mathcal{A}^{\Theta_1}$ and $f_2 \in \mathcal{A}^{\Theta_2}$, where $\Theta_1, \Theta_2 \subset \Theta$. The belief f_1 stochastically dominates (in inverse of the reverse hazard rate) the belief f_2 on $\Theta_1 \cap \Theta_2$ if $F_1(\theta)/f_1(\theta) \leq F_2(\theta)/f_2(\theta)$ for all $\theta \in \Theta_1 \cap \Theta_2$.

Lemma 1. *Let $f_1 \in \mathcal{A}^{\Theta_1}$ and $f_2 \in \mathcal{A}^{\Theta_2}$, where $\Theta_1, \Theta_2 \subset \Theta$, be such that f_1 stochastically dominates the belief f_2 on $\Theta_1 \cap \Theta_2$. Then*

$$\bar{x}^{f_1}(\theta) > \bar{x}^{f_2}(\theta) \quad \text{and} \quad \bar{U}_a^{f_1}(\theta) > \bar{U}_a^{f_2}(\theta) \tag{13}$$

for all $\theta \in \Theta_1 \cap \Theta_2$.

The finding that the optimal decision $\bar{x}^f$ is decreasing in the rate F/f will be the crux of our welfare results. Lemma 1 implies that using the described dominance concept the agent can rank some admissible beliefs if they have the same support. But a similar preference relation over the beliefs is not available for the society (or the principal). In other words, on a given support of positive length there exists no belief of the regulator which is desired most by the whole society. However, this negative result is not disappointing for us. Indeed, as the rest of this paper will make it clear, there are situations where the social welfare is very sensitive to the support of beliefs that are believed to contain the searched type parameter.

Hereafter, we fix and denote by θ^T the private type parameter of the agent, and define $\Theta^T = \{\theta^T\}$. Now we consider a single-stage learning prior to regulation, which changes the current information structure (f^0, Θ^0) to (f^1, Θ^1) where $f^i \in \mathcal{A}^{\Theta^i}$ and $\Theta^1 \subset \Theta^0$ with $\Theta_1 \notin \{\Theta^0, \Theta^T\}$. We further suppose that the regulator has not acquired any additional information about the distribution of the types in the finer support Θ^1. Then the posterior belief f^1 on Θ^1 should be obtained by some (pre-announced) update rule from the prior f^0 on Θ^0.

Here we simply assume that the learning of the regulator is exogenous, and moreover the underlying learning technology is such that it always pays to spend on learning from the viewpoint of the society. In the following definition we state the minimal restriction on f^1 to ensure that the information structure (f^1, Θ^1) is superior to (f^0, Θ^0).

Definition 2. The structure (f^1, Θ^1) contains *valuable* (more) information about θ^T than the structure (f^0, Θ^0) if $\Theta^1 \subset \Theta^0$ and $f^1(\theta^T)/f^0(\theta^T) \geq f^1(\theta)/f^0(\theta)$ for all $\theta \in \Theta^1$.

In the single-stage learning we consider the information about θ^T is incomplete. Thus, more information resulting from learning does not necessarily imply that the regulator, and the society, are aware of its presence. Indeed, one can naturally ask the following question: can the regulator be ever *certain* that he has "more information" under some incomplete learning? Note that the regulator can simply check whether Θ^1 is a subset of Θ^0. So, the above question boils down to whether the regulator can certify that $f^1(\theta^T)/f^0(\theta^T) \geq f^1(\theta)/f^0(\theta)$ for all $\theta \in \Theta^1$ without actually knowing what the value of θ^T is. Apparently, the answer is 'yes' only if $f^1(\theta)/f^0(\theta)$ is constant over Θ^1. This observation leads us to focus on the following belief update rule.

Definition 3. The belief f^1 on Θ^1 is the Bayesian update of f^0 on Θ^0 where $\Theta^1 \subset \Theta^0$ if $f^1(\theta) = f^0(\theta)(1+\gamma)$ for all $\theta \in \Theta^1$, where $\gamma = [\int_{\Theta^1} f(\theta)d\theta]^{-1} - 1$.

Then, a regulator can convince the society that he knows more about the regulated agent only if the regulator is a Bayesian learner. We state this result, which requires no further proof, as follows:

Proposition 2. *The regulator knows that the structure* (f^1, Θ^1) *contains more information about* θ^T *than the structure* (f^0, Θ^0) *only if* f^1 *is the Bayesian update of* f^0.

In sequel, we point to situations in which the agent prefers the Bayesian regulator to have more information about his private type.

Proposition 3. *Suppose the regulator knows that the learned structure* (f^1, Θ^1) *contains more information than the prior structure* (f^0, Θ^0)*, where* $\min(\Theta^1) > \min(\Theta^0)$ *and* $\max(\Theta^1) = \max(\Theta^0)$. *Then the welfare of the regulated agent is higher under the learned structure, i.e.* $\bar{U}_a^{f^1}(\theta) > \bar{U}_a^{f^0}(\theta)$ *for all* $\theta \in \Theta^1$.

With Bayesian learning that shrinks the type space from the left, the regulator's posterior belief stochastically dominates his prior belief. Then the welfare of the agent increases by Lemma 1, whereas the changes in the welfare of the principal and the society are ambiguous. The below corollary to Proposition 3 points to the potential of honest signalling of the agent about his type space before the implementation of the regulatory mechanism.

Corollary 1. *Let* (f^0, Θ^0) *be the current information structure and the regulator be known to use Bayes rule in updating his beliefs. Then the agent finds it profitable to signal that his type parameter cannot be in the interval* $[\min(\Theta^0), \theta^T)$.

The following proposition symmetrically examines learning with right-sided contraction of the type space.

Proposition 4. *Suppose the regulator knows that the learned structure* (f^1, Θ^1) *contains more information than the prior structure* (f^0, Θ^0)*, where* $\min(\Theta^1) = \min(\Theta^0)$ *and* $\max(\Theta^1) < \max(\Theta^0)$. *Then the welfare of the regulated agent is lower whereas the welfare of the principal and the society are both higher under the learned structure, i.e.* $\bar{U}_a^{f^1}(\theta) < \bar{U}_a^{f^0}(\theta)$,

$\bar{U}_p^{f^1}(\theta) > \bar{U}_p^{f^0}(\theta)$ *and* $\bar{W}^{f_1}(\theta) > \bar{W}^{f^0}(\theta)$ *for all* $\theta \in \Theta^1$.

Note that Bayesian learning that shrinks the type space only from the right leaves the inverse of the reverse hazard rate, hence the optimal decision variable, unchanged. Nevertheless, the informational rents of the agent become reduced as the upper bound of the integral expression in (8) becomes smaller under the new information structure. With lower informational rents, the social welfare in (9) becomes higher independently from the weight α of the agent's welfare. It follows that the welfare of the principal, which coincides with the social welfare when $\alpha = 0$, becomes higher, too. Obviously, the regulator must keep on this kind of learning until a point where the expected gain of getting more information is balanced by the cost of learning.

4. Conclusions

In a generalized principle-agent model, we have examined a Bayesian regulator's learning about the private information of the regulated agent. We have specified what 'more information' means and demonstrated that more information about the informed agent needs not be undesirable for him. We have also characterized situtations in which the principal and the society benefit from the regulator's learning.

Our findings support the view that one should be careful in determining what to expect from Bayesian mechanisms with their existing specifications. It has long been noticed that the subjective nature of beliefs may cast some doubts on the implementability of the Bayesian mechanisms. Crew and Kleindorfer (1986), Vogelsang (1988), Koray and Sertel (1990) criticized the Bayesian approach in regulation on the grounds of unaccountability and manipulability of the regulator's subjective prior beliefs. In a very recent study, Koray and Saglam (2005) examine the same issue in the Baron and Myerson (1982) model of monopoly regulation. They show that all interest groups in the society are extremely sensitive to the prior belief of the regulator. There exist beliefs yielding values arbitrarily close to the supremum of actual welfare and expected welfare of the regulated agent (monopolist) and the principal (consumers), respectively. Moreover, under some other beliefs one can come as close to the infimum of actual welfare of both parties as possible. When the belief of the regulator is unverifiable by the public, the existence of such critical beliefs leads to a bargaining game over the beliefs between a corrupt or captured regulator and the interest groups in the society, which distorts the regulatory outcome predicted by Baron and Myerson (1982).

What we add to the previous results is that Bayesian mechanisms may yield unpredictable and sometimes undesirable outcomes even in the presence of a benevolent and sincere regulator if the socially efficient type of learning is not completely specified as part of the regulatory mechanism.

References

Baron, D., and R.B. Myerson (1982) "Regulating a Monopolist with Unknown Costs" *Econometrica* **50**, 911-930.

Blume, L., and D. Easley (1995) "What has the Rational Learning Literature Taught Us?" in *Learning and Rationality in Economics* by A. Kirman and M. Salmon, Eds., Oxford: Blackwell, 12-39.

Cox, J.C., Shachat J., and M. Walker (2001) "An Experimental Test of Bayesian Learning in Games" *Games and Economic Behavior* **34**, 11-33.

Crew, M.A., and P.R. Kleindorfer (1986) *The Economics of Public Utility Regulation,* Cambridge, MA: MIT Press.

Gibbard, A. (1973) "Manipulation of Voting Schemes: A General Result" *Econometrica* **41**, 587-602.

Guesnerie, R., and J.J. Laffont (1984) "A Complete Solution to a Class of Principal-Agent Problems with an Application to the Control of a Self-Managed Firm" *Journal of Public Economics* **25**, 329-369.

Jordan, J.S. (1991) "Bayesian Learning in Normal Form Games" *Games and Economic Behavior* **3**, 60-81.

Kalai, E., and E. Lehrer (1993) "Rational Learning Leads to Nash Equilibria" *Econometrica* **61**, 1019-1045.

Koray, S., and I. Saglam (2005) "The Need for Regulating a Bayesian Regulator" *Journal of Regulatory Economics* **28**, 5-21.

Koray, S., and M.R. Sertel (1990) "Pretend-but-Perform Regulation and Limit Pricing" *European Journal of Political Economy* **6**, 451-472.

Myerson, R.B. (1979) "Incentive Compatibility and the Bargaining Problem" *Econometrica* **47**, 61-74.

Vogelsang, I. (1988) "A Little Paradox in the Design of Regulatory Mechanisms" *International Economic Review* **29**, 467-476.

Appendix

Proof of Proposition 1. The integrand in the objective function of (11) is differentiated with respect to $x(\theta)$ to obtain the optimality condition (12). Using the asssumptions (A2) and (A7), it is easy to check that the same integrand is concave in x.

To show that the solution to (12) satisfies the incentive-compatibility constraint (IC), we will first prove that the optimal solution $\bar{x}$ is nonincreasing in θ. Total differentiation of (12) with respect to θ yields

$$\left(\frac{\partial^2 V_p}{\partial x^2} + \frac{\partial^2 V_a}{\partial x^2} + (1-\alpha)\frac{F(\theta)}{f(\theta)}\frac{\partial^3 V_a}{\partial^2 x^2 \partial\theta}\right)\frac{d\bar{x}}{d\theta} =$$

$$\left(-(1-\alpha)\frac{d}{d\theta}\left(\frac{F(\theta)}{f(\theta)}\right) - 1\right)\frac{\partial^2 V_a}{\partial x\partial\theta} - \frac{\partial^2 V_p}{\partial x\partial\theta} - (1-\alpha)\frac{F(\theta)}{f(\theta)}\frac{\partial^3 V_a}{\partial x\partial\theta^2}.$$

Using the assumptions (A2), (A3), (A4), (A6) and (A7) together with the assumption that $F(\theta)/f(\theta)$ is nondecreasing in θ, we conclude that $d\bar{x}/d\theta$ is nonpositive.

The net utility of the agent when he truthfully reports his type as θ is

$$U_a(\theta) = -\int_{\theta}^{\theta_1} \frac{\partial}{\partial\tilde{\theta}} V_a(\bar{x}(\tilde{\theta}),\tilde{\theta})d\tilde{\theta}$$

by (8). The net utility of the agent when he misreports its unknown parameter as $\hat{\theta}$ while θ is the true parameter is

$$U_a(\theta,\hat{\theta}) = V_a(\bar{x}(\hat{\theta}),\theta) + U_a(\hat{\theta}) - V_a(\bar{x}(\hat{\theta}),\hat{\theta}). \tag{14}$$

Subtracting $U_a(\theta)$ from (14) we get

$$\begin{aligned} U_a(\theta,\hat{\theta}) - U_a(\theta) &= -\int_{\hat{\theta}}^{\theta} \frac{\partial}{\partial\tilde{\theta}} V_a(\bar{x}(\tilde{\theta}),\tilde{\theta})d\tilde{\theta} + V_a(\bar{x}(\hat{\theta}),\theta) - V_a(\bar{x}(\hat{\theta}),\hat{\theta}) \\ &= -\int_{\hat{\theta}}^{\theta} \frac{\partial}{\partial\tilde{\theta}}\Big(V_a(\bar{x}(\tilde{\theta}),\tilde{\theta}) - V_a(\bar{x}(\hat{\theta}),\tilde{\theta})\Big)d\tilde{\theta} \leq 0 \end{aligned}$$

from (A4) and $d\bar{x}(\theta)/d\theta \leq 0$. Thus, the optimal program (12) is incentive-compatible.

Finally to check condition (IR), i.e. $U_a(\theta) \geq 0$ at the optimal solution $\bar{x}$, is straightforward from (8) thanks to assumption (A5). ∎

Proof of Lemma 1. Total differentiation of (12) at the optimal decision $\bar{x}^f$ with respect to $F(\theta)/f(\theta)$ yields

$$\left(\frac{\partial^2 V_p}{\partial x^2} + \frac{\partial^2 V_a}{\partial x^2} + (1-\alpha)\frac{F(\theta)}{f(\theta)}\frac{\partial^3 V_a}{\partial^2 x^2 \partial\theta}\right)\frac{d\bar{x}^f}{d[F(\theta)/f(\theta)]} = -(1-\alpha)\frac{\partial^2 V_a}{\partial x\partial\theta}.$$

From assumptions (A2), (A4) with strict inequality and (A7) it follows that $\bar{x}^f$ is decreasing in $F(\theta)/f(\theta)$. Considering equation (8), using assumptions (A4) and (A5) and

$F_1(\theta)/f_1(\theta) < F_2(\theta)/f_2(\theta)$, we conclude that $\bar{U}_a^{f_1}(\theta) > \bar{U}_a^{f_2}(\theta)$ for all $\theta \in \Theta$. ∎

Proof of Proposition 3. Since f^1 is a Bayesian update of f^0 on a finer support, $f^1(\theta) > f^0(\theta)$ and hence $F^1(\theta) < F^0(\theta)$ for all $\theta \in [\min(\Theta^1), \max(\Theta^0))$ while $F^1(\max(\Theta^0)) = F^0(\max(\Theta^0)) = 1$. This implies that $F^1(\theta)/f^1(\theta) < F^0(\theta)/f^0(\theta)$ for all $\theta \in \Theta^1$. Then from Lemma 1, $\bar{x}^{f^1}(\theta) > \bar{x}^{f^0}(\theta)$ and $\bar{U}_a^{f^1}(\theta) > \bar{U}_a^{f^0}(\theta)$ for all $\theta \in \Theta^1$. ∎

Proof of Proposition 4. Since f^1 is a Bayesian update of f^0, $f^1(\theta) = f^0(\theta)(1+\gamma)$ for all $\theta \in \Theta^1$, where $\gamma = [F(\max(\Theta_1))]^{-1} - 1$. Note that $F^1(\theta)/f^1(\theta) = F^0(\theta)/f^0(\theta)$ and therefore $x^{f^1}(\theta) = x^{f^0}(\theta)$ for all $\theta \in \Theta^1$. Then from (8) we obtain $\bar{U}_a^{f^1}(\theta) < \bar{U}_a^{f^0}(\theta)$, since $\max \Theta^1 < \max \Theta^0$.

We have $\bar{W}^{f^1}(\theta) > \bar{W}^{f^0}(\theta)$ since $\bar{W}^{f^1}(\theta) = \bar{V}_p^{f^1} + \bar{V}_a^{f^1} - (1-\alpha)\bar{U}_a^{f^1} = \bar{W}^{f^0}(\theta) + (1-\alpha)(\bar{U}_a^{f^0} - \bar{U}_a^{f^1})$. Finally, $\bar{U}_p^{f^1}(\theta) > \bar{U}_p^{f^0}(\theta)$ follows from the fact that $\bar{W}^{f^0}(\theta) = \bar{U}_p^{f^0}(\theta)$ when $\alpha = 0$. ∎