



City Research Online

City, University of London Institutional Repository

Citation: Ungureanu, S. (2013). Dynamic Contracting under Permanent and Transitory Private Information (13/07). London, UK: Department of Economics, City University London.

This is the unspecified version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/2921/>

Link to published version:

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.



School of Social Sciences
CITY UNIVERSITY LONDON

Department of Economics

Dynamic Contracting under Permanent and Transitory Private Information

Sergiu Ungureanu
City University London

Department of Economics
Discussion Paper Series
No. 13/07

Dynamic Contracting under Permanent and Transitory Private Information

Sergiu Ungureanu*

November 5, 2013

Abstract

To understand how firms create and maintain long term relationships with consumers, or how procurement relations evolve over time, it is useful to study a dynamic variant of the classical two-type-buyer contract in mechanism design. It is less trivial and more interesting if the utility determinant (or utility type) is not fixed or completely random, and fair assumptions are that it is either stochastic, or given by a distribution whose parameters are common knowledge. The first approach is that of Battaglini (2005), while the second is pursued in this paper. With two possible types of buyers, the buyer more likely to have a high utility type will receive the first-best allocations, while the other will receive the first best only if he has the high utility type.

Keywords: dynamic contracting, mechanism design, truthful reporting, information structure, learning.

1 Introduction and Related Literature

The relationships between buyers and sellers are often dynamic in nature, and the relevant private information will rarely stay hidden when there is repeated interaction. This has important implications for the pricing of products or the transfer of goods between a principal and an agent, since the inefficiency induced by asymmetric information may be reduced by considering the long term behavior of the agent.

There are many ways to model repeated interactions between, say, a buyer and a seller, and by the revelation principle the optimal mechanism will involve the

*Department of Economics, City University London, EC1V 0HB, UK, scu2@duke.edu. I want to thank my doctoral advisor, Prof. Curtis Taylor, and the seminar participants at Duke University, for comments, suggestions and essential help in directing my work.

seller reporting his private information, as long as adequate incentive compatibility and rationality conditions are satisfied. But, if the private information known before signing the contract has no bearing on future private information, the problem is in a sense trivial. The dynamic contract becomes a pasting of static ones. It is more interesting to consider that there is a link between the private information at different times, and this link can be stochastic, or statistic – given by an underlining condition like a probability distribution, or in general a combination. In Battaglini (2005), we see the simplest form of a stochastic link: the utility type evolves according to a Markov process. Here I considers the simplest statistic link: the utility types are given by a distribution depending on a buyer type. The buyer type could be common knowledge, but this would make the problem trivial again. I will consider chiefly the buyer's type to be the buyer's private information, but the problem becomes even more interesting once we explore the idea of a hidden or partially hidden buyer type, which can be learned by both parties through repeated interactions.

It is important to look at such dynamic contractual relationships, since they are increasingly important for the economy. As information technology advances, tracking customers becomes easier and can lead to better tailored products and offers. Moreover, handling such long term data can improve the customer screening potential of companies which have the resources. A great example is the market for mobile telephony. Mobile phone carriers have the ability to enforce long term contracts, and the service usage can be carefully monitored. Moreover, the service to users can be charged at different rates based on quantity brackets, and can package utility increasing features into the contract – phone quality, text messaging, data plans. The contracts themselves come with initial or overtime financial commitments. The consumer can arguably have knowledge of his potential usage pattern, defining a possible buyer type, and his monthly consumption choice would be equivalent to reporting a utility type by the taxation principle. Consider also a monopolistic supplier selling to a small business or a company dealing with its franchise. The company or franchise hold specific private information about the profitability of the business, as well as about the day to day sales potential, which would determine short term profits. Long term strategy must incorporate the information from repeated interaction.

This paper presents the interaction between a buyer and a seller, with the buyer has private information about his "utility type" and his "buyer type". The buyer type will determine the utility type statistically, but all buyer types can have one of the same utility types, so there is some overlap which makes buyer types unseparable in a trivial manner. Considering the buyer type as private information, it is clear that the buyer's type will influence the type of contract that is optimally offered only to him, as well as in a setting in which there are more than one buyer type. It turns out that there is separation between buyer types, which for two types happens through basically two contracts, where for reports of a low utility (type) the quantities allocated differ.

An alternative approach to a similar problem in current literature is in Battaglini (2005). He considers a dynamic interaction between a buyer and a seller much like

in this paper, where utility only can change from a utility type determinant, and this type is linked across time by a Markov process. This assumption may describe the setting of a monopolistic supplier or a franchise owner better, at least when it can be said that the profitability of the agent's business is subject to a Markov process because, say, the daily number of returning customers is Markov. In other situations it makes more sense to consider that the profitability of the business is simply determined by unchangeable factors – here the agent type, and randomness.

In a working paper, Pavan et al. (2009) attempt to characterize general dynamic mechanism design problems in finite time, and again for an infinite horizon in Pavan et al. (2010). They consider a general stochastic process for information, with possibly non-time-separable agent payoffs. To obtain more explicit results, they require continuity in the probability measure that determines types, and in general, continuity in the total variation metric for the most general results. This does not translate to the finite type case, where the differential conditions cannot apply. One important observation they make is that the vanishing distortions at the bottom principle (VDB) in Battaglini (2005) is not a robust result. In general, distortions need not be monotonic in type or in time, nor vanish in the long run.

Courti and Li (2000) study a somewhat similar problem with two stages, where the first stage is where the agent, an airline ticket buyer, finds out his type, which restricts somewhat the buyer's possible valuations for the ticket. Considering contracts with partial refunds, the authors find that the informativeness of the signal the buyer gets is what determines the optimal mechanisms, and not the uncertainty that affects all the buyer types.

2 Describing the Model

Consider a setting with a monopolistic principal, or seller, and an agent, or buyer. The buyer's period by period utility function is given by

$$U(\theta, q) = \theta q - q^2/2 - p(q),$$

where q is the buyer's consumption, p is a price he pays for it henceforth determined by q , and θ is a utility, or demand, type determinant which can take the values $\{\theta_h, \theta_l\}$. The buyer has ex-ante private information about the probability of having a demand of type θ_h instead of θ_l in any one period. This probability is determined by the buyer's type, and can take the values $\{\alpha_h, 1 - \alpha_h\}$ or $\{\alpha_l, 1 - \alpha_l\}$. There is common knowledge that the only two buyer types appear with odds $\{\varphi, 1 - \varphi\}$ respectively. Later on, this assumption will be relaxed and generalized. The reservation utility for the buyer is 0, and the seller has the power to commit to a contract. Both buyer and seller are risk neutral and have no liquidity constraints, so they care only about maximizing expected utility. We can restrict attention to direct truthful mechanisms. The game starts in period 0, with the buyer reporting his α value, after which he may have to make a payment, although we may assume *w.l.o.g.* that the payment is deferred. In subsequent rounds, the buyer reports his θ




FIGURE 1. Moves by nature. The vertical dotted line shows when the contract is signed.

demand determinant, receives an allocation q and makes a new payment – Figure 1. Let generic $\theta_h > \theta_l \geq 0$, $1 > \alpha_h > \alpha_l > 0$, $\varphi \in (0, 1)$, but strict inequalities are only to consider the interesting problem. In the following, I will focus on a game with infinite periods, but sometimes I may also consider a finite period game for description or as a stepping stone. Unless explicitly stated, the results will hold for the finite as well as infinite settings. In the finite case I will ignore discounting.

3 Solving the Model

It is interesting to consider both the finite and the infinite time periods versions. For the infinite setting, we could possibly use a recursive approach, but simple considerations show that the number of relevant state variables grows linearly in time, and it is not immediate in what way one can summarize the relevant history. Therefore, looking at the optimal direct mechanism for the finite game is easier, and could possibly be extended with a continuity argument to the infinite setting. For that, one typically considers rationality and compatibility incentive constraints, which are then incorporated in a global maximization problem. However, as the number of time periods grows, the number of incentive compatibility conditions will also grow polynomially. It is important then to par down this number to the absolute minimum, before proceeding to any analysis in this direction.

3.1 Reducing the number of ICC conditions

The first step in the solving strategy is proving that the set of incentive compatibility conditions can be restricted somewhat. This claim will be proven for a more general buyer utility function, and an arbitrary finite number of utility types l . The utility function is $U_t(\theta, q_t(\cdot))$ in round t , discounted if needed, where θ can take any value from the finite set $\{\theta_1, \dots, \theta_l\}$, and the allocation $q_t(h_0, h_1, \dots, h_t)$ is a function of the history of messages the buyer sent. Fix $U_0 \equiv 0$, since we can delay payment. Here α is just an index for one of a finite number of probability distributions over the values θ , and it has full support *w.l.o.g.* We can consider the messages h_t to be determined by a function of the buyer's true type and demand type history, $h^t \equiv \sigma_t(\alpha, \theta^1, \theta^2, \dots, \theta^t)$, $h^0 \equiv \sigma_0(\alpha)$, where $\sigma \equiv \{\sigma_t(\cdot)\}$ is the player's strategy. Any direct truthful mechanism that implements a menu will have to satisfy individual rationality constraints (IRCs) for each α , since the buyer has ex-ante information about his type before he signs a contract.

$$\forall \alpha : \quad 0 \leq \sum_{t=1}^n \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(\alpha, \theta^1, \dots, \theta^t)),$$

where $\Theta^t = (\theta^1, \dots, \theta^t)$ is one possible history of types for period t , and $\lambda_\alpha(\theta^i)$ is the probability of getting the value θ^i that is given by the distribution indicated by α . For clarification, in our two-types model $\alpha_h \equiv \lambda_{\alpha_h}(\theta_h)$, $1 - \alpha_h \equiv \lambda_{\alpha_h}(\theta_l)$, $\alpha_l \equiv \lambda_{\alpha_l}(\theta_h)$, and $1 - \alpha_l \equiv \lambda_{\alpha_l}(\theta_l)$. The incentive compatibility constraints (ICCs) for a direct truthful mechanism are

$$\begin{aligned} \forall \sigma, \alpha : \quad & \sum_{t=1}^n \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(\alpha, \theta^1, \dots, \theta^t)) \geq \\ & \sum_{t=1}^n \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, h^1, \dots, h^t)). \end{aligned} \quad (1)$$

Proposition 3.1. *The set of incentive compatibility conditions (1) is satisfied if a stronger set of compatibility conditions (\mathcal{S}) hold, whereby in each period the buyer considers only one-time deviations from truth, followed by truthful reporting in the future. That is, the ICCs are implied by the following set of inequalities, for¹ $k \in \overline{0, n}$, $\forall \sigma, \alpha, \Theta^k$:*

$$\begin{aligned} & U_k(\theta^k, q_k(h^0, h^1, \dots, h^{k-1}, \theta^k)) + \\ & + \sum_{t=k+1}^n \sum_{\Theta^t | \Theta^k} \prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(h^0, h^1, \dots, h^{k-1}, \theta^k, \theta^{k+1}, \dots, \theta^t)) \geq \\ & \geq U_k(\theta^k, q_k(h^0, h^1, \dots, h^k)) + \\ & + \sum_{t=k+1}^n \sum_{\Theta^t | \Theta^k} \prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(h^0, h^1, \dots, h^k, \theta^{k+1}, \dots, \theta^t)). \end{aligned}$$

¹To be exact, consider that $\theta^0 = \alpha$. At $k = 0$ there are only the incentive compatibility conditions for each type α , which are the same as the ICCs.

Proof. For now let $n < \infty$. The proof is by induction. First consider a finite setting with n periods. Let σ be any strategy, and consider a deviation from this strategy to truth telling, only in the last period n , on only one of the possible histories. Assume the following inequality holds. $\forall \sigma, \alpha, \Theta^n$:

$$U_n(\theta^n, q_n(h^0, h^1, \dots, h^{n-1}, \theta^n)) \geq U_n(\theta^n, q_n(h^0, h^1, \dots, h^{n-1}, h^n)). \quad (2)$$

It is just a special case of (1), and of (\mathcal{S}) when $k = n$. Now, for a generic σ , it implies that $\forall \alpha$:

$$\begin{aligned} & \sum_{t=1}^n \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, h^1, \dots, h^t)) \leq \\ & \sum_{t=1}^{n-1} \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, h^1, \dots, h^t)) + \sum_{\Theta^n} \prod_{i=1}^n \lambda_\alpha(\theta^i) U_n(\theta^n, q_n(h^0, h^1, \dots, h^{n-1}, \theta^n)), \end{aligned}$$

so it makes sense to always report truthfully in the last period. This is the first step in the induction process. In general, consider that the following holds for an arbitrary $0 \leq k \leq n$, i.e. that it is better to report truthfully from period $k+1$ on. $\forall \sigma, \alpha$:

$$\begin{aligned} & \sum_{t=1}^n \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, h^1, \dots, h^t)) \leq \\ & \sum_{t=1}^k \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, h^1, \dots, h^t)) + \\ & + \sum_{t=k+1}^n \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, h^1, \dots, h^k, \theta^{k+1}, \dots, \theta^t)). \end{aligned}$$

Assume that \mathcal{S} hold. \mathcal{S} say that, assuming future truthful reporting, the buyer finds the total discounted utility from reporting his true type today is larger than for any other report. Notice that this has to hold for all α , and for any history prior to period k . Using \mathcal{S} , rework the previous expression. $\forall \sigma, \alpha$:

$$\begin{aligned} & \sum_{t=1}^n \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, h^1, \dots, h^t)) \leq \\ & \sum_{t=1}^{k-1} \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, h^1, \dots, h^t)) + \sum_{\Theta^k} \prod_{i=1}^k \lambda_\alpha(\theta^i) U_k(\theta^k, q_k(h^0, h^1, \dots, h^k)) + \\ & \sum_{t=k+1}^n \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, h^1, \dots, h^k, \theta^{k+1}, \dots, \theta^t)) = \\ & = \sum_{t=1}^{k-1} \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, h^1, \dots, h^t)) + \sum_{\Theta^k} \prod_{i=1}^k \lambda_\alpha(\theta^i) \left[U_k(\theta^k, q_k(h^0, h^1, \dots, h^k)) + \right. \\ & \left. \sum_{t=k+1}^n \sum_{\Theta^t | \Theta^k} \prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(h^0, h^1, \dots, h^k, \theta^{k+1}, \dots, \theta^t)) \right] \leq \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{t=1}^{k-1} \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, h^1, \dots, h^t)) + \sum_{\Theta^k} \prod_{i=1}^k \lambda_\alpha(\theta^i) \left[U_k(\theta^k, q_k(h^0, h^1, \dots, h^{k-1}, \theta^k)) + \right. \\
&\quad \left. \sum_{t=k+1}^n \sum_{\Theta^t | \Theta^k} \prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(h^0, h^1, \dots, h^{k-1}, \theta^k, \dots, \theta^t)) \right] = \\
&\quad \sum_{t=1}^{k-1} \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, h^1, \dots, h^t)) + \\
&\quad + \sum_{t=k}^n \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, h^1, \dots, h^{k-1}, \theta^k, \dots, \theta^t)).
\end{aligned}$$

This result is the induction step, which takes us to $k = 1$. The last step is applying the 0-period ICC, which is also in \mathcal{S} . Therefore, we have used expressions \mathcal{S} to obtain the ICCs.

Now consider $n = \infty$. The partial sum from k to ∞ , of the series that gives the buyer's utility is

$$\sum_{t=k}^{\infty} \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, h^1, \dots, h^t)),$$

which is bounded and must go to 0 as $k \rightarrow \infty$, if the function U_t is exponentially bounded because of discounting. For an infinite direct mechanism that implements a menu, say that the infinite set of constraints \mathcal{S} will hold. Let σ represent the optimal strategy, mapping any buyer type and sequence of utility types to optimal reports. From this, devise an approximate strategy that is a "truncated" version of σ , in which the buyer reports according with the strategy σ up to period k , and then he reports truthfully. Then it must be that the payoff from the new strategy is an increasingly good approximation as k increases, since only the partial sum from k on is changed. Fix k , and apply the compatibility conditions in \mathcal{S} from period k backwards to prove, again by induction, that the completely truthful strategy is weakly better than the "truncated" σ , which approximates in payoff the original. As k goes to infinity, we obtain that truthful reporting is weakly optimal in the context of a menu satisfying \mathcal{S} , which simply means that the regular ICCs have to hold too. \square

Intuitively, the alternative conditions of Proposition 3.1 allow us to improve, step by step, any non-truthful strategy by considering deviations towards truthful reporting. Starting with an arbitrary strategy σ the last period, it must be that, if it is preferable for the buyer to not lie regardless of the past history, he will improve on σ by reporting truthfully. This must hold after any history, so we immediately have that in the last period \mathcal{S} implies truthful reporting. This allows us to move to the previous period, where we again observe, history by history, that truthful reporting is preferred, and so on by backward induction. This is possible since there are compatibility conditions for any α type at each node; i.e., it is optimal for the buyer not to lie, assuming truthful reporting in the future, even if he isn't

supposed to get this menu because he misrepresented his α . It is thus possible to improve onto any σ by deviating towards the truthful strategy.

The result holds for $n = \infty$ also, and the proof is by contradiction. Starting again with a possibly non-truthful strategy σ , we can approximate it by assuming that after period k the buyer decides to report truthfully. Only after period k will his utility be changed by the new strategy, and the change is small if k is large. This is true because we will need to incorporate discounting into U_t , and if the period by period utility is bounded, then the partial sum over time that gives the buyer's total utility must converge. With the new strategy it is possible to use conditions in \mathcal{S} as before, to prove by backward induction that true reporting is superior to the approximated σ . Then, in the limit $k \rightarrow \infty$, we must have that the true strategy is as good as σ . But this statement is precisely the requirement that truthful strategies satisfy the ICCs, which is what we need to prove.

Proposition 3.2. *The set of ICCs in (1) imply a weaker set of compatibility conditions (\mathcal{W}), in which the buyer, who has reported his true type α if in periods after 0th, considers only one-time deviations from truth, followed by truthful reporting in the future. That is, the ICCs imply the following inequalities, for² $k \in \overline{0, n}$, $\forall \sigma, \alpha, \Theta^k$:*

$$\begin{aligned}
& U_k(\theta^k, q_k(\alpha, h^1, \dots, h^{k-1}, \theta^k)) + \\
& \quad + \sum_{t=k+1}^n \sum_{\Theta^t | \Theta^k} \prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(\alpha, h^1, \dots, h^{k-1}, \theta^k, \theta^{k+1}, \dots, \theta^t)) \geq \\
& \geq U_k(\theta^k, q_k(\alpha, h^1, \dots, h^k)) + \\
& \quad + \sum_{t=k+1}^n \sum_{\Theta^t | \Theta^k} \prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(\alpha, h^1, \dots, h^k, \theta^{k+1}, \dots, \theta^t)), \text{ and} \\
& \sum_{t=1}^n \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(\alpha, \theta^1, \dots, \theta^t)) \geq \sum_{t=1}^n \sum_{\Theta^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, \theta^1, \dots, \theta^t)).
\end{aligned}$$

Proof. The idea is to observe that the compatibility conditions expressed in \mathcal{W} can be seen as just a subset of the ICCs. For the last inequality above, it is obvious. In general, consider the ICC for two strategies, the truthful strategy and a deviation to the strategy that reports truthfully except after a history $(\alpha, h^1, \dots, h^{k-1})$, when $\sigma_k(\alpha, h^1, h^2, \dots, h^{k-1}, \theta^k) = h^k$ and, subsequently, again truthful reporting after any continuation history. Define $H^k \equiv (\alpha, h^1, \dots, h^{k-1}, \theta^k)$. You can simplify the ICC by eliminating the payoffs for the histories $\Theta^n - \Theta^n | H^k$ and their continuations,

²At $k = 0$ consider $\theta^0 = \alpha$, and $U_0 = 0$ as defined.

which must be equal, so you get:

$$\begin{aligned}
& \prod_{i=1}^k \lambda_\alpha(\theta^i) U_k(\theta^k, q_k(\alpha, h^1, \dots, h^{k-1}, \theta^k)) + \\
& \quad + \prod_{i=1}^k \lambda_\alpha(\theta^i) \sum_{t=k+1}^n \sum_{\Theta^t | H^k} \prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(\alpha, h^1, \dots, h^{k-1}, \theta^k, \theta^{k+1}, \dots, \theta^t)) \geq \\
& \geq \prod_{i=1}^k \lambda_\alpha(\theta^i) U_k(\theta^k, q_k(\alpha, h^1, \dots, h^k)) + \\
& \quad + \prod_{i=1}^k \lambda_\alpha(\theta^i) \sum_{t=k+1}^n \sum_{\Theta^t | H^k} \prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(\alpha, h^1, \dots, h^k, \theta^{k+1}, \dots, \theta^t))
\end{aligned}$$

Simplify the common factor to get the generic compatibility condition of \mathcal{W} . In this construction, we haven't assumed $n < \infty$, so the result holds for the infinite setting. \square

The intuition behind the proof rests on the observation that, if reporting truthfully is always better than not, then it is also better than lying only once, and that the continuation utilities for a person who considers lying once are the same as for a person who may have misrepresented his past utility types – θ , as long as he is the same buyer type α . This is true because future payoffs depend only on reports and the buyer type, which are both assumed the same. So having compatibility conditions in each period checking whether a buyer prefers lying for that period to true reporting is equivalent to considering some of the ICCs.

We can therefore say $\mathcal{W} \subset \text{ICCs}$, and $\mathcal{W} \subset \mathcal{S}$ is by definition. Proposition 3.1 says that if a menu satisfies \mathcal{S} , then it also satisfies the ICCs, so we can think of the set of possible menus under \mathcal{S} as more restrictive than those that satisfy the ICCs. In general, a menu implemented by a mechanism will be a point in the menu parameter space. The sets of conditions $\mathcal{W}, \mathcal{S}, \text{ICCs}$ will delimit a subset of this space, where the menus inside obey the conditions. Let $\mathcal{D}[\cdot]$ denote the subset. Then we must have $\mathcal{D}[\mathcal{S}] \subset \mathcal{D}[\text{ICCs}] \subset \mathcal{D}[\mathcal{W}]$ (Figure 2). Now consider the maximization problem that determines the optimal mechanism. To find the optimal truthful mechanism we need to restrict the domain of menus to $\mathcal{D}[\text{ICCs}]$. A containing domain may generate a false maximum, if it is located outside $\mathcal{D}[\text{ICCs}]$, while a subdomain can lead to maxima on its boundaries that could be suboptimal. To solve for the optimal mechanism, my strategy will be to find the optimal menu in $\mathcal{D}[\mathcal{W}]$ and show that it is in $\mathcal{D}[\mathcal{S}]$.

3.2 Fixing some of the allocations

The next step is to reduce the possible ways in which a menu representing a truthful mechanism can be constructed. From now on, we return to the two-type setting, although the results will hold for an arbitrary number of types, if properly stated. This section shows through a few lemmas that the optimal menu must have first

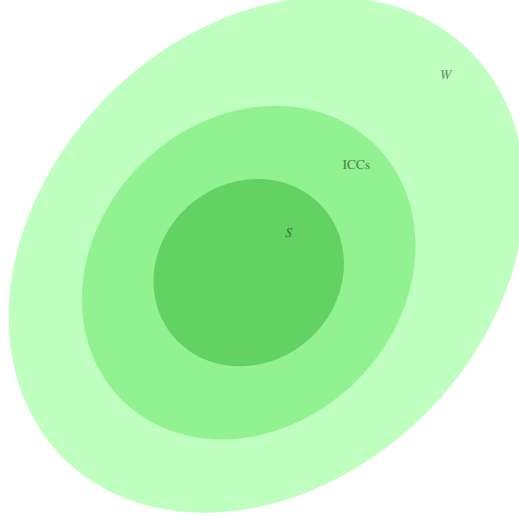


FIGURE 2. The restricted domains according to $\mathcal{W}, \mathcal{S}, ICCs$ in the menu parameter space.

best allocations for reports of θ_h , as well as for θ_l after type α_h was reported. The lemmas work for any $n \leq \infty$, and also if we rephrase them with \mathcal{S} instead of \mathcal{W} .

Lemma 3.1. *Consider any period with allocations. Denote with q_h, q_l the quantities and p_h, p_l the prices in the menu offered to the buyer in that period, contingent upon the θ -type report, and let u_h, u_l be the expected continuation utilities for the buyer, given by the buyer type and truthful reporting in the future. Then, if conditions \mathcal{W} hold, it must be that $q_h \geq q_l$.*

Proof. Assume the \mathcal{W} compatibility conditions in the period hold. By them it is *w.l.o.g.* that in the future the player will report truthfully. Then there are only two continuation values that are needed to describe a player's incentives, say u_h and u_l , and:

$$\begin{aligned} \theta_h q_h - q_h^2/2 - p_h + u_h &\geq \theta_h q_l - q_l^2/2 - p_l + u_l \\ \theta_l q_l - q_l^2/2 - p_l + u_l &\geq \theta_l q_h - q_h^2/2 - p_h + u_h. \end{aligned}$$

We can rewrite the equations to:

$$\begin{aligned} \theta_h(q_h - q_l) &\geq p_h - p_l + q_h^2/2 - q_l^2/2 + u_l - u_h \\ p_h - p_l + q_h^2/2 - q_l^2/2 + u_l - u_h &\geq \theta_l(q_h - q_l), \end{aligned}$$

from which you can deduce that $(\theta_h - \theta_l)(q_h - q_l) \geq 0$. Since $\theta_h > \theta_l$, we must have that $q_h \geq q_l$. This lemma remains true if we replace \mathcal{W} with \mathcal{S} , because we would have similar inequalities for each pair u_l^α, u_h^α , leading to the same conclusion. \square

The result can be understood to mean that, if $q_l > q_h$ and if a buyer with θ_l prefers the low allocation and price to the high ones, then the buyer with θ_h must

prefer it too, because his total utility is increased more by an increase in the allocation, and future utility is determined only by today's report, and not by today's true type. But this evidently contradicts a compatibility condition.

Lemma 3.2. *Consider any period with allocations in the game described. Denote with q_h, q_l the quantities and p_h, p_l the prices in the menu offered to the buyer in that period, and let u_h, u_l be the expected continuation utilities. Assuming that \mathcal{W} hold, it must be that $q_h \geq \theta_h$ and $q_l \leq \theta_l$ in the optimal menu.*

Proof. Prove by contradiction, i.e., assume that $q_h < \theta_h$. Then $\exists \epsilon > 0$ s.t. $q_h + \epsilon < \theta_h$. Now consider the following changes in the menu offered: $q_h \rightarrow q_h + \epsilon$, and $p_h \rightarrow p_h + (\theta_h - q_h)\epsilon - \epsilon^2/2$. Observe that both q_h and p_h are strictly increased.

The new expected utility of the buyer, when he truthfully reports θ_h , is:

$$\theta_h(q_h + \epsilon) - (q_h + \epsilon)^2/2 - (p_h + (\theta_h - q_h)\epsilon - \epsilon^2/2) + u_h = \theta_h q_h - q_h^2/2 - p_h + u_h.$$

When the buyer reports θ_l , everything stays the same. Now verify that the new incentive compatibility conditions will also be satisfied:

$$\begin{aligned} \theta_h(q_h + \epsilon) - (q_h + \epsilon)^2/2 - (p_h + (\theta_h - q_h)\epsilon - \epsilon^2/2) + u_h &\geq \theta_h q_l - q_l^2/2 - p_l + u_l \\ \theta_l q_l - q_l^2/2 - p_l + u_l &\geq \theta_l(q_h + \epsilon) - (q_h + \epsilon)^2/2 - (p_h + (\theta_h - q_h)\epsilon - \epsilon^2/2) + u_h \\ &\Leftrightarrow \end{aligned}$$

$$\begin{aligned} \theta_h q_h - q_h^2/2 - p_h + u_h &\geq \theta_h q_l - q_l^2/2 - p_l + u_l \\ \theta_l q_l - q_l^2/2 - p_l + u_l &\geq \theta_l q_h - q_h^2/2 - p_h - (\theta_h - \theta_l)\epsilon + u_h. \end{aligned}$$

Because $(\theta_h - \theta_l)\epsilon > 0$, the old incentive compatibility conditions imply that these hold. Any other constraints, either incentive compatibility constraints at other nodes or individual rationality constraints, will not be changed, because the expected utility for the buyer does not change. Therefore we have found that there is a way to strictly improve the seller's expected return, without changing the expected payoff of the buyer. So this allocation is not optimal, therefore $q_h \geq \theta_h$. For the statement $q_l \leq \theta_l$, prove by contradiction by assuming that $q_l \geq \theta_l$. Then $\exists \epsilon > 0$ s.t. $q_l - \epsilon > \theta_l$. Now change the menu offered: $q_l \rightarrow q_l - \epsilon$, and $p_l \rightarrow p_l + (q_l - \theta_l)\epsilon - \epsilon^2/2$. Observe that q_l is strictly decreased while p_l is strictly increased. The new expected utility for the buyer that truthfully reports θ_l does not change:

$$\theta_l(q_l - \epsilon) - (q_l - \epsilon)^2/2 - (p_l + (q_l - \theta_l)\epsilon - \epsilon^2/2) + u_l = \theta_l q_l - q_l^2/2 - p_l + u_l.$$

Neither does the utility of the buyer that truthfully reports θ_h . The new incentive compatibility constraints are:

$$\begin{aligned} \theta_h q_h - q_h^2/2 - p_h + u_h &\geq \theta_h(q_l - \epsilon) - (q_l - \epsilon)^2/2 - (p_l + (q_l - \theta_l)\epsilon - \epsilon^2/2) + u_l \\ \theta_l(q_l - \epsilon) - (q_l - \epsilon)^2/2 - (p_l + (q_l - \theta_l)\epsilon - \epsilon^2/2) + u_l &\geq \theta_l q_h - q_h^2/2 - p_h + u_h \end{aligned}$$

\Leftrightarrow

$$\begin{aligned}\theta_h q_h - q_h^2/2 - p_h + u_h &\geq \theta_h q_l - q_l^2/2 - p_l - (\theta_h - \theta_l)\epsilon + u_l \\ \theta_l q_l - q_l^2/2 - p_l + u_l &\geq \theta_l q_h - q_h^2/2 - p_h + u_h.\end{aligned}$$

These new conditions must hold if the old ones did. Also, no other compatibility or rationality conditions will be changed. So this change improves the seller's return, therefore it must be that $q_l \geq \theta_l$ was suboptimal. \square

This lemma shows that the θ_h type must always be allocated at least the first best in an optimal mechanism, because you can always increase a suboptimal allocation, and then charge a higher price to compensate his utility increase, and this change will not make the report θ_h better for any lower utility type. Similarly, allocating more than the first best to the low utility type is suboptimal because, if the allocations is greater, then the seller can decrease it marginally, and compensate by an increase in price, and this compensation would leave the lowest type indifferent, but for all the other types it would be insufficient, therefore leaving any constraints relaxed.

Lemma 3.3. *Assume that in the setting presented \mathcal{W} hold, and the results of the previous lemmas. Then, for every period, for each history, at most one of the compatibility conditions in \mathcal{W} for an agent α will bind.*

Proof. Write again the ICCs:

$$\begin{aligned}\theta_h q_h - q_h^2/2 - p_h + u_h &\geq \theta_h q_l - q_l^2/2 - p_l + u_l \\ \theta_l q_l - q_l^2/2 - p_l + u_l &\geq \theta_l q_h - q_h^2/2 - p_h + u_h.\end{aligned}$$

Assume that both conditions above hold with equality. Then

$$\begin{aligned}\theta_h(q_h - q_l) &= q_h^2/2 + p_h - u_h - q_l^2/2 - p_l + u_l \\ \theta_l(q_h - q_l) &= q_h^2/2 + p_h - u_h - q_l^2/2 - p_l + u_l,\end{aligned}$$

so $\theta_h(q_h - q_l) = \theta_l(q_h - q_l)$. Because $q_h \geq \theta_h > \theta_l \geq q_l$, we have a contradiction. \square

From the previous lemma, we can also see that the high and low utility type allocations must be different, by the conditions in \mathcal{W} . If more than one compatibility conditions hold, then it says that the difference in utilities minus the difference in prices for two alternatives has to be 0, and for more than one utility type. This is not possible because utility differences for the same allocations are strictly monotone in θ . A more general formulation of utility that would preserve the result has $\mathcal{U}_\theta(\theta, q) > 0$, and $\mathcal{U}_{\theta, q}(\theta, q) > 0$ (known as the *Spence-Mirrlees single-crossing condition*).

Lemma 3.4. *Consider any period with allocations, and assume the compatibility condition in \mathcal{W} for the buyer with θ_h -type utility does not bind. Then $q_l \geq \theta_l$. Similarly, if the θ_l -type condition does not bind, then $q_h \leq \theta_h$.*

Proof. Say that $q_l < \theta_l$. Then $\exists \epsilon > 0$ s.t. $q_l + \epsilon < \theta_l$. Consider the menu transformation $q_l \rightarrow q_l + \epsilon, p_l \rightarrow p_l + (\theta_l - q_l)\epsilon - \epsilon^2/2$. Then, the new incentive compatibility conditions will be:

$$\begin{aligned} \theta_h q_h - q_h^2/2 - p_h + u_h &\geq \theta_h(q_l + \epsilon) - (q_l + \epsilon)^2/2 - p_l + (\theta_l - q_l)\epsilon - \epsilon^2/2 + u_l \\ \theta_l(q_l + \epsilon) - (q_l + \epsilon)^2/2 - p_l + (\theta_l - q_l)\epsilon - \epsilon^2/2 + u_l &\geq \theta_l q_h - q_h^2/2 - p_h + u_h. \end{aligned}$$

$$\Leftrightarrow$$

$$\begin{aligned} \theta_h q_h - q_h^2/2 - p_h + u_h &\geq \theta_h q_l - q_l^2/2 - p_l + u_l + (\theta_h - \theta_l)\epsilon \\ \theta_l q_l - q_l^2/2 - p_l + u_l &\geq \theta_l q_h - q_h^2/2 - p_h + u_h. \end{aligned}$$

As long as ϵ is small enough, the first incentive compatibility condition will hold for the new menu. Moreover, the expected revenue of the seller is strictly increased without affecting the buyer's utility. For the second part of the statement, consider now $q_h \rightarrow q_h - \epsilon, p_h \rightarrow p_h + (q_h - \theta_h)\epsilon - \epsilon^2/2$. After reworking, the new incentive compatibility conditions will be:

$$\begin{cases} \theta_h q_h - q_h^2/2 - p_h + u_h \geq \theta_h q_l - q_l^2/2 - p_l + u_l \\ \theta_l q_l - q_l^2/2 - p_l + u_l \geq \theta_l q_h - q_h^2/2 - p_h + u_h + (\theta_h - \theta_l)\epsilon. \end{cases}$$

If the θ_l compatibility condition will not bind, then for a small enough ϵ the above conditions will hold. As before, the seller's expected revenue is increased, but the buyer's payoffs are left the same. \square

Decreasing the high utility allocations marginally towards the first best and then compensating the price will leave the high utility type indifferent, but it will improve the high report option for the other type. However, if the old compatibility constraint was slack, a marginal increase in the high type allocation will not violate the it. Similarly, you can argue that increasing the low type allocation towards the first best and then compensating the low utility type with a higher price will also be an improvement, but only if the high type's compatibility constraint is slack.

3.3 Solving the maximization problem

From now on assume \mathcal{W} hold and that in every period with allocations the θ_l -type incentive compatibility condition does not bind, while for the case with $\sigma_0 = \alpha_h$ the θ_h -type compatibility condition will also not have to bind, i.e., it is slack. With the previous lemmas, the assumption gives us that, in every period after the 0th, $q_h = \theta_h$, and that after $\sigma_0 = \alpha_h$ we must have in every subsequent period $q_l = \theta_l$. Moreover, at period 0, we assume that only the rationality constraint for the α_l -type and the compatibility constraint for the α_h -type bind. The proof requires that we verify the assumptions after we find the solution.

With the assumptions above, the allocations after reporting α_h are fixed. The seller needs to optimize the low demand type allocations if the buyer reported α_l in period 0. Since we have assumed that only the α_l -type compatibility condition binds in period 0, we can maximize the seller's value, keeping track of the fact that the information rent earned by a α_h buyer pretending to have the α_l -type must be subtracted with the appropriate weight.

Lemma 3.5. *The simplified seller's maximization problem is given by the following recursive form:*

$$P(v) = \max_{\{p_h, p_l, v_h, v_l, q_l\}} (1 - \varphi)[\alpha_l p_h + (1 - \alpha_l) p_l] + \alpha_l \delta P(v_h) + (1 - \alpha_l) \delta P(v_l) + \varphi[-\alpha_h(\theta_h^2/2 - p_h) - (1 - \alpha_h)(\theta_l q_l - q_l^2/2 - p_l)] \quad (3)$$

$$\begin{aligned} \text{s.t. (i)} \quad & \theta_h^2/2 - p_h + \delta v_h = \theta_h q_l - q_l^2/2 - p_l + \delta v_l, \\ \text{(ii)} \quad & v = \alpha_l(\theta_h^2/2 - p_h) + (1 - \alpha_l)(\theta_l q_l - q_l^2/2 - p_l) + \alpha_l \delta v_h + (1 - \alpha_l) \delta v_l, \\ \text{(iii)} \quad & q_l \geq 0. \end{aligned}$$

Here v is the utility value for the buyer at the beginning of period 1, v_h, v_l are his continuation utilities if he reports θ_h and θ_l respectively. Since we assumed that the α_l -type rationality constraint in period 0 holds with equality, we must have that $v = 0$ in the first period. Because there is an indeterminacy in the price allocation, we can also fix $v_h = v_l = 0$ to find explicit prices, but after we have determined the quantity allocations from the maximization problem.

Proposition 3.3. *Assuming that (iii) doesn't bind, the allocation after reporting θ_l in the maximization problem (3) is always given by*

$$q_l = \theta_l - (\theta_h - \theta_l) \frac{\varphi}{1 - \varphi} \cdot \frac{\alpha_h - \alpha_l}{1 - \alpha_l}$$

and with the normalization $v_h = v_l = 0$,

$$\begin{aligned} p_h &= \frac{\theta_h^2}{2} - (\theta_h - \theta_l)(1 - \alpha_l)q_l, \\ p_l &= (\alpha_l(\theta_h - \theta_l) + \theta_l)q_l - \frac{q_l^2}{2}. \end{aligned}$$

(iii) holds with equality when $\frac{\theta_l(1 - \alpha_l)}{\theta_h(\alpha_h - \alpha_l) + \theta_l(1 - \alpha_h)} \leq \varphi \leq 1$, and in this case $q_l = 0$, $p_h = \frac{\theta_h^2}{2}$ and $p_l = 0$ in every period.

Proof. First assume that (iii) doesn't bind. Solve (i) and (ii) for p_h and p_l :

$$\begin{aligned} p_h &= -v + \delta v_h + \frac{\theta_h^2}{2} - q_l(1 - \alpha_l)(\theta_h - \theta_l), \\ p_l &= -\frac{q_l^2}{2} - v + v_l \delta + q_l(\alpha_l(\theta_h - \theta_l) + \theta_l). \end{aligned}$$

Replace these expressions into the maximization problem, and get the first order conditions. The first order condition determining q_l will be:

$$\frac{(1 - \alpha_l)(-1 + \varphi)q_l}{-1 + \varphi} + \frac{\theta_l(1 + \alpha_h\varphi - \varphi - \alpha_l) - (\alpha_h - \alpha_l)\theta_h\varphi}{-1 + \varphi} = 0.$$

This solves for the q_l expression. Observe that the allocation is determined for every period after period 0. Because we ignored the non-negativity constraint (iii), it is necessary to see if q_l is positive. It turns out that a necessary and sufficient condition for that is:

$$\varphi \leq \frac{\theta_l(1 - \alpha_l)}{\theta_h(\alpha_h - \alpha_l) + \theta_l(1 - \alpha_h)}.$$

When this condition doesn't hold, then $q_l = 0$. To find the prices, it is necessary to fix the indeterminacy by making some assumptions, because the seller can commit on the contract offered, therefore changing the all the prices at period t by ϵ and compensating at period $t + k$ by $\delta^k\epsilon$ leaves the buyer with equivalent choices. Moreover, there is a one-to-one trade-off for the buyer and the truthful seller between v_h or v_l and p_h, p_l . This suggests that neither the absolute, nor the relative values of the continuation utilities v_h, v_l are fixed, and the easiest way to deal with them is to assume $v_h = v_l = 0$. Of course, we have to check that the compatibility constraints of the buyer that reports non-truthfully will also hold, but this assumption makes all the four compatibility constraints for each node hold, because of (i). Using the assumptions, we get the prices p_h, p_l in every period. \square

The next result gives the allocations and prices offered if the buyer reports α_h in period 0.

Proposition 3.4. *The allocations after reporting α_h are $q_h = \theta_h$ and $q_l = \theta_l$, and*

$$p_h = \frac{\theta_h^2}{2} - \theta_h\theta_l + \alpha_h\theta_h\theta_l + \theta_l^2 - \alpha_h\theta_l^2$$

$$p_l = \alpha_h\theta_h\theta_l + \frac{\theta_l^2}{2} - \alpha_h\theta_l^2,$$

as well as a transfer of $\frac{(\alpha_h - \alpha_l)(\theta_h - \theta_l)q_l}{1 - \delta}$ from the seller to cover the information rent for the α_h type buyer.

Proof. Solve the following equations, representing the binding θ_h -type incentive compatibility and the buyer value identity respectively, after assuming again $v = v_h = v_l = 0$:

$$\alpha_h \left(\frac{\theta_h^2}{2} - p_h \right) + (1 - \alpha_h) \left(\frac{\theta_l^2}{2} - p_l \right) = 0,$$

$$\frac{\theta_h^2}{2} - p_h = \theta_h\theta_l - \frac{\theta_l^2}{2} - p_l,$$

and this gives the results. We have assumed that the θ_h -type compatibility condition must not bind, so that we can get $q_l = \theta_l$, yet we have obtained the above solution making the exact opposite assumption. But that is fine, since the compatibility condition is slack. To see that, change the solution above with a small value κ :

$$p_h = \frac{\theta_h^2}{2} - \theta_h\theta_l + \alpha_h\theta_h\theta_l + \theta_l^2 - \alpha_h\theta_l^2 - (1 - \alpha_h)\kappa$$

$$p_l = \alpha_h\theta_h\theta_l + \frac{\theta_l^2}{2} - \alpha_h\theta_l^2 + \alpha_h\kappa.$$

It is easy to check that, because the high and low demand compatibility conditions cannot both bind at the same time, a κ small enough always exists that will leave every condition satisfied with strict inequality. This allows us to deduce that $q_l = \theta_l$ was a good assumption. \square

It is left is to check that all the assumptions on which individual rationality or compatibility conditions in \mathcal{W} will bind have been correct. Because of the symmetry of the allocations after reporting α , it is obvious that all the compatibility conditions in \mathcal{S} are also satisfied. It was assumed that the θ_h -type condition holds with equality, for the buyer who has reported his true α , which means that the θ_l -type compatibility condition will also hold. If the buyer has not reported his α -type truthfully, then the compatibility conditions starting from period 1 will be essentially unchanged, because of the symmetry of the allocations, which means that $v_h = v_l$.

The last thing we need to check are the conditions in period 0. The IRCs are satisfied by construction. If an α_h -type buyer reports α_l , then the information rent he earns is given by:

$$v^h = (\alpha_h - \alpha_l)(\theta_h - \theta_l)q_l + \delta\alpha_h v^h + \delta(1 - \alpha_h)v^h,$$

which solves to

$$\frac{(\alpha_h - \alpha_l)(\theta_h - \theta_l)q_l}{1 - \delta}.$$

This is the transfer a α_h -type buyer will get in period 0 in our solution, and the expected utility starting from period 1 will be 0 by construction. Similarly, we can find the expected utility of an α_l -type buyer who reports α_h :

$$\frac{(\alpha_h - \alpha_l)(\theta_h - \theta_l)(q_l - \theta_l)}{1 - \delta} < 0.$$

Therefore, the lower type wouldn't report α_h , so he gets a continuation utility of 0. Because we have found a solution satisfying the stronger collection of compatibility conditions in \mathcal{S} , it is the correct solution for the set of ICCs.

4 Discussion

From the solution in Proposition 3.3, it can be seen intuitively that the distortion of θ_l allocations is induced by the need to decrease rents for the α_h type buyer. A marginal increase Δq in all allocations for θ_l reports of a α_l buyer, at any period, would lead to a proportional increase in rents $\varphi(\alpha_h - \alpha_l)(\theta_h - \theta_l)\Delta q$, and the extra revenue for the seller is $(1 - \varphi)(1 - \alpha_l)(\theta_l - q_l)\Delta q$, in the first order approximation. At equilibrium, the two must be equal, and this gives the result.

Now vary the different parameters, to see if the results coincide with those of simpler problems. If $\varphi = 0$, then there is only an α_l type, who must sign a contract before he knows his first θ . As expected, our solution becomes $q_l = \theta_l, q_h = \theta_h$, which is the classical dynamic contract when the buyer has no private information. If $\varphi \rightarrow 1$, this leads to $q_l \rightarrow 0$, and to 0 rent for the high type buyer, because the α_h buyer becomes increasingly important, so his rent is weighted progressively more. Again, in the limit we find the classical dynamic contract with one buyer type and no private information. The same holds true when $\alpha_h = \alpha_l$, and it is because we don't have to distinctive buyer types anymore.

It is interesting to contrast the result with those of Battaglini (2005), who considers that there is only one buyer type, but his utility determinant θ evolves according to a Markov process, from a known prior distribution, and assumes that high types will be correlated. He showed that the optimal dynamic contract for such a setting converges to first best allocations as time passes. Furthermore, only reporting a sequence of θ_l would lead to suboptimal allocations; once the buyer has reported θ_h , it would be followed only by first best allocations. In this setting, allocation distortions are permanent for a low buyer type, which leads to consistent losses of social welfare. This comes in part from the fact that a Markov process becomes progressively less connected to past information, so the buyer's private information – which usually leads to the allocation distortions – is in a sense weaker. However, it is also arguable that the two settings are fundamentally different, because in Battaglini (2005) the contract is signed after the buyer knows his first θ value. We can also modify our problem to consider such a situation (pending work), and the result is that there is a difference only in that first period, when if θ_l is reported then $q = \theta_l - (\theta_h - \theta_l)Pr[\theta_h]/Pr[\theta_l]$, after which allocations revert to the ones here. It is also interesting to note that, if we consider only two periods of the game mentioned, then we must have the same solution as Battaglini, since the information process is trivially Markov, and this checks out.

Consider now the problem faced by a mobile phone carrier. He must distinguish between two types of buyers, one that is either relatively poor, or that has a lifestyle that doesn't require many phone calls, and one who is likely to consume many minutes. By the taxation principle, the contracts offered to them are the implementation of a truthful mechanism. The result presented here suggests that the optimal choice of contracts would have basically two, for each buyer type. Let's say that one contract has sufficient minutes included, and one has a flat pay-per-use. When the buyer has a high utility draw, it means that his need for the call is great, so he will take the call no matter what his phone plan is. If his utility

draw is low, then the flat rate buyer will decide, maybe, to not make the call, or talk less; but the buyer of the plan with minutes need not bother. It is also obvious that the buyer who expects to have many calls will prefer the better plan, and he will have to pay less per minute. The difference can be considered his information rent, because he had the option to choose the flat rate plan in the beginning. This description is consistent with typical consumer behavior. One can also think of contracts that may or may not include additional features, like messages or data plans, where the utility draw is over the usefulness of the feature. People who believe they will often need such features are more likely to buy a contract that will allow them their use more often.

5 Conclusion

This paper deals with long term contracts between a principal and an agent, and considers that the agent has private information that is not completely random or fixed, but comes from the same prior distribution, which is his private information. The optimal contracts for two types of buyers with two types of utilities have some of the usual properties from the static problem, like no distortions for high utility draws, but are radically different from the static or the infinite games with random or persistent information. There is consistent distortion for the low type buyer, when he reports a low utility, and the distortion decreases social welfare. The results also provide some insight into the pricing strategies of monopolistic suppliers of firms or franchises, and of mobile phone carriers.

References

- [1] Athey, S. and Segal, I., "An Efficient Dynamic Mechanism," *Draft*, February 2007.
- [2] Battaglini, M., "Long-Term Contracting with Markovian Consumers," *The American Economic Review*, June 2005, Vol. 85, No.3, 637–658.
- [3] Bergeman, D. and Valimaki, J., "The Dynamic Pivot Mechanism," *Econometrica*, March 2010, Vol. 78, No. 2 , 771–789.
- [4] Bolton, P. and Dewatripont, M., "Contract theory," Cambridge, MA: MIT Press, 2005.
- [5] Courty, P. and Li, H., "Sequential Screening," *Review of Economic Studies*, 2000, 67, 697–717.
- [6] Miravete, E., "Choosing the Wrong Calling Plan? Ignorance and Learning," *The American Economic Review*, 2003, 93(1), 297–310.
- [7] Pavan, A., Segal, I., Rhodes, B., Starmer, C. and Sugden, R., "A Test of the Theory of Reference-Dependent Preferences," *The Quarterly Journal of Economics*, 1997, 2, 479–505.
- [8] Pavan, A., Segal, I. and Toikka, J., "Dynamic Mechanism Design: A Myersonian Approach," *Working Paper*, August 9 th, 2013.
- [9] Pavan, A., Segal, I. and Toikka, J., "Infinite-Horizon Mechanism Design: The Independent Shock Approach," *Working Paper*, July 16th, 2010.