



City Research Online

City, University of London Institutional Repository

Citation: Ungureanu, S. (2017). Dynamic Contracting under Permanent and Transitory Private Information. *International Journal of Game Theory*, 46, pp. 667-692. doi: 10.1007/s00182-016-0552-8

This is the published 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/11860/>

Link to published version: <https://doi.org/10.1007/s00182-016-0552-8>

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.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

Dynamic contracting under permanent and transitory private information

Sergiu Ungureanu¹

Accepted: 25 September 2016 / Published online: 6 October 2016
© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract I examine a model of long-term contracting in which the buyer is privately informed about the discrete probability distribution for his future value for a divisible product, and fully characterize the optimal long term contract that will be offered by a monopolistic seller in a simple case where two types of buyers can have two types of utility in any period. In such a case, the buyer more likely to have a high utility type will receive the first-best allocations indifferent of his value report, while the lower type will receive the first best only if he makes a high utility report. The paper also supplements the current literature on infinite dynamic games with continuous buyer types, which relies on the use of a distribution of types with full support and an envelope theorem. With discrete types, the number of compatibility constraints considered can be greatly reduced by sandwiching the border of the space of solutions allowed by constraints: formulate the maximization problem in a wider space with fewer constraints and prove that the solution obeys a simpler set of stronger constraints that places it in the allowed region.

Keywords Asymmetric information · Dynamic contracting · Mechanism design · Information structure

JEL Classification C73 · D82

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.

✉ Sergiu Ungureanu
Sergiu.Ungureanu.1@city.ac.uk

¹ Department of Economics, City University London, Northampton Square, London EC1V 0HB, UK

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 a principal and an agent, say a seller and a buyer, and by the revelation principle the optimal mechanism will involve the buyer reporting his private information, as long as adequate incentive compatibility and rationality conditions are satisfied. Because of truthful reporting, if the private information known when the contract is signed is the only thing that has a bearing on future private information the problem can be seen intuitively in a greatly simplified way. The dynamic contract resembles a string of static ones, contingent on the time period, agent type and period specific reports.

An early contribution in the literature on dynamic mechanisms with adverse selection is the work of [Baron and Besanko \(1984\)](#), who describe a monopoly able to commit, dealing with an agent of variable type, in two periods. Dynamic settings with constant types are considered in [Dewatripont \(1989\)](#), [Hart and Tirole \(1988\)](#), and [Laffont and Tirole \(1988\)](#), [Laffont and Tirole \(1990\)](#). [Cremer and McLean \(1988\)](#) justifies the usual assumption of independent information across agent types to prevent full rent extraction. [Laffont and Tirole \(1990\)](#) introduce a standard framework, analyze the two-period renegotiation proof contract, and [Laffont and Tirole \(1996\)](#) introduce variable agent types and commitment in the context of regulation. More dynamic models in which types are stochastically determined are studied by [Rustichini and Wolinsky \(1995\)](#), and [Courty and Li \(2000\)](#). In more recent contributions, [Battaglini \(2005\)](#) studies variable stochastic types in infinite time with commitment, and [Battaglini \(2007\)](#) analyzes limited commitment with stochastic types. [Esó and Szentes \(2007\)](#) consider many agents in two time periods, when the principal can also release private information. [Bergemann and Välimäki \(2010\)](#) and [Athey and Segal \(2013\)](#) introduce dynamic generalizations of the classic Vickery-Clarke-Groves incentive compatible and efficient mechanism, [Boleslavsky and Said \(2013\)](#) study an infinite period dynamic interaction with persistent shocks and continuous types, while [Pavan et al. \(2014\)](#) work on the general mechanism design problem with continuous types.

From an application perspective, it is interesting to consider that there is a link between the private information at different times, and this link can be stochastic—given by a probabilistic evolution; statistic—given by an underlining condition like an initial 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. It is simplest in the sense that the information is linked only between consecutive periods, and the preceding periods to the last are irrelevant. By contrast, this paper considers a statistic link: the utility types are given by a distribution depending on a buyer type.

This paper presents the interaction between a buyer and a seller, when the buyer has private information about his periodic *utility type*, and his long run *buyer type*. The

buyer type will be informative of the utility type in terms of probabilities, and all buyer types will draw from the same discrete set of utility types. The buyer type is private information and will influence the type of contract that is optimally offered only to him. A general model for the interaction between a seller and a buyer with one of a discrete set of types is developed. General results allowing for the simplification of the problem are obtained, based on the idea of generating a restricted set of constraints and finding a candidate mechanism in the larger parameter space restricted by that subset, and then verifying that a sufficiently strong set of constraints also hold. The results are applied to derive expressions for optimal allocations in a setting in which each period the buyer can have one of two possible utility types, a low type and a high type, which can change over time. The utility types are selected according to a distribution which can also change, with the condition that the high buyer type is more likely to have the higher utility type each period. The analysis shows that there is separation between buyer types: for the two types, two contracts are offered, and for each of these contracts, reporting different utility types in a period leads to different allocations in that period. In a typical way, allocations for unprofitable low buyer types have to be reduced, to reduce the information rent paid to the profitable high types. Therefore, allocations are optimal following high utility reports of any buyer types, and also optimal for low utility reports of the high buyer types, but are distorted for low buyer types reporting low utility types. Prices in the optimal menu can be chosen to compensate buyers for an optimal level of information rent. As expected, the level of distortion is increasing in the relative weight of high types, and increasing in the difference of the effects of private information. If the buyer type is reported in a truthful mechanism in the first period, any dynamic changes of allocations are a reflection of time changes of the possible type values, which are common knowledge. The contribution of this worked application is to extend the analysis of the dynamic interaction between a principal and an agent of two possible types, who has private information about his future demand and about the likelihood of that demand being high or low, also allowing the demand types and the likelihood to vary in a general way over time.

It is important to look at such dynamic contractual relationships, since they are of major importance in an economy. As information technology advances, tracking customers becomes easier and can lead to better tailored products and offers (Miravete 2003). Handling such long term data can improve the customer screening potential of companies which have the resources. Examples are businesses that sell memberships to repeat clients. They 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 the seller can package utility increasing features into the contract. The contracts themselves come with initial or overtime financial commitments. The consumer will arguably have knowledge of his potential usage pattern, defining a possible buyer type, and his monthly consumption choice would be equivalent to reporting a period specific type—by the taxation principle. One can consider also monopolistic suppliers selling to a small business, or franchise owners dealing with individual franchises. The agent companies and franchises 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.

The approach of this paper is novel also because of the method of analysis employed. As is standard in the literature, incentive compatibility constraints necessary for a direct truthful mechanism have to be introduced in the maximization problem. The goal is to restrict the relevant set of incentive compatibility constraints.¹ To simplify the problem consider, in a first stage, consider only the compatibility constraints that have to be satisfied after the contract is signed, and after the buyer reports truthfully the private information he has *ex ante*. These constraints ignore the buyers that have misreported their initial private information. With only this subset of constraints, say \mathcal{W} , the optimal menu space considered is wider than allowed, so any menu found satisfying \mathcal{W} that is also incentive compatible *ex ante* will be fully optimal. With the problem thus divided, the next step involves finding conditions that restrict the large space of viable parameters—these conditions are equivalent to \mathcal{W} holding. Then maximization conditions are applied to this restricted set of parameters, so that the solution is fully incentive compatible. Because distortions in allocations arise as a result of *ex ante* private information, the more analytically involved problem is figuring out correct allocations that satisfy the incentive compatibility constraints between *ex ante* types. The benefit of this method is that the analytically difficult part can be separated, and expressions are solved in a problem with far fewer parameters.

This type of contracting problem is important since agent private information that comes in sequentially is realistic in many settings. As a qualitative justification for this assertion, consider the problem faced by a gym, trying to sell gym memberships. For simplicity, assume there are only two kinds of buyers: a heavy gym user type, and a light user type. The light user type is more likely to only want access to basic machines and facilities, while the heavy user type will more often require specific facilities (pool, squash court, fitness classes), in addition to the basics. In this example, the heavy and light users are our high and low types, and the momentary need for facilities in the gym is the period by period utility type. Both the heavy and the light users know what they will likely want, but do not yet know when that would happen. The difference is understood to be in the frequency of this demand. The results suggests a low type should be offered a basic membership package, with pay-per-use access to the facilities not included, while a high type should be offered an extensive package, which comes with a discount on more facilities, and which is worthwhile only for a heavy user. The discount would correspond to the information rent paid to the high type, since he has the option of buying the basic package.

Courty and Li (2000) study a similar problem about the allocation of an indivisible good, where private information is linked across only two time periods. Both a discrete two-type case and a continuum of types are considered. In the first period the agent finds out his type, which partially restricts the buyer's possible valuations for the ticket in the future. The informativeness of the buyer's private information can differ both in the sense of first order stochastic dominance—when the high type knows he values a ticket more than the low type, or in the sense of a mean-preserving spread—when the high type is more uncertain about the value of the ticket. The first period information can be used by the principal to propose an interim contract to increase his returns,

¹ In settings with continuous types, additional assumptions on the distribution related to the inverse hazard rate or first order stochastic dominance have precisely this function.

compared to signing the contract *ex post* when the agent has all the private information. 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. The virtual value determined in the continuous type case is suggestive of the allocations determined in this paper. In both cases, the distortion of the low type is proportional to the informativeness of the initial agent type and the relative difference in the utility types. The private information that the buyer has before signing the contract is what leads to information rents and to distortions of future allocations from the first-best. Therefore, one would expect that the present set-up induces distortions at arbitrary periods similarly, to account for the same private information. However, that cannot be directly deduced, since the utility type reports in future periods can provide information on the initial buyer type, so the dynamic game is not separable into independent games at different periods. The present paper can be seen as extending some of their results to an arbitrary, possibly infinite, number of periods.

[Battaglini \(2005\)](#) considers an alternative approach, with a dynamic interaction between a buyer and a seller in infinite time, much like in this paper, where agent utility can change period by period based on a discrete utility type determinant, and this type is linked across time by a Markov (memoryless) process. It is assumed that types are positively correlated across time, that is, a high type is more likely than a low type to become a high type again. The incipient type evolution in the first two periods is the same as in the present paper, which will consider that types are correlated across time based on a common distribution, so it is no surprise that the allocations are also the same. After the second period, the dynamics diverge: in [Battaglini \(2005\)](#) the distortion for the low report type remains only after a history of low-type reports, and is probabilistically reduced progressively until it vanishes at infinity; in this paper the distorted allocation for a low type report is maintained indifferent of the history of low-type reports, and does not converge to the first best. [Rustichini and Wolinsky \(1995\)](#) work with another simpler model with Markov types in which the strategic agent behavior is limited, and the focus is on the principal's learning of the demand environment.

[Boleslavsky and Said \(2013\)](#) examine a model with a dynamic interaction between a buyer and a seller, which shares many of the features of the model herein. The significant differences are that (i) they use a continuous buyer distribution with positive support—i.e., there are no gaps in buyer types, and that (ii) the shocks to the period by period utility type of the buyer have a very specific structure, with persistent effect—the buyer's value is proportional to the product of the history of shocks. The no-gap-in-buyer-type assumption is necessary, given the analysis method, much like in [Pavan et al. \(2014\)](#), because of the use of the envelope theorem to restrict incentive compatibility constraints. The assumption of the buyer value as a product of shocks is however more consequential. It makes any shock to the valuation of the buyer infinitely persistent, and its multiplicative form leads to a “path-independence” property for allocations, which is used in reducing the number of compatibility constraints considered, thus simplifying the analysis. The optimal allocations derived start with “honeymoon phases” of advantageous prices of different lengths, and demanding different entry costs, which serve to screen the buyer types.

In a very important recent paper, Pavan et al. (2014) attempt to characterize general dynamic mechanism design problems in finite and infinite time. 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. The results do not translate to discrete type settings, where the differential methods cannot apply, however, they are quite general, and the intuition described is informative. In particular, the authors note that the change in informativeness of the period-by-period report of the buyer determines the change in time of the allocation distortion. In the present paper, the optimal truthful mechanism requires reporting the buyer type in the first period, which implies that the allocation distortions in subsequent periods should be the same, since future reports are equally informative for the seller.

2 A general discrete model for the buyer's problem

Consider a setting with a monopolistic principal, or seller, and an agent, or buyer. The seller and the buyer are interacting over possibly infinitely many periods $t \in \{0, 1, \dots, n\}$. The buyer can be of one of a finite set of types, indexed by α , and the ex-ante distribution of α types is common knowledge. Furthermore, let θ be the period-by-period utility type of the buyer, which can take any value from the finite set $\{\theta_1, \dots, \theta_l\}$, and define the probability mass function of buyer type α over utility types θ to be $\lambda_\alpha(\cdot)$, which has full support.

Let the buyer's utility function at time t be $U_t(\theta^t, q_t(\cdot), p_t(\cdot))$, possibly time dependent, discounted if needed.² For any α and t , θ^t can take any of the finite set of values $\{\theta_1, \dots, \theta_l\}$, and the allocation $q_t(h^0, h^1, \dots, h^t)$ and price $p_t(h^0, h^1, \dots, h^t)$ are a function of the history of messages h^0, h^1, \dots, h^t previously sent by the buyer. At time 0, the buyer is assumed to report his buyer type α , receive no allocation and make no payment. Because the seller will be able to commit, he can delay payment so $U_0 \equiv 0$ is *w.l.o.g.*, with the outside option of the buyer normalized to 0. The reports h^t are determined by a function of the buyer's true type, demand type history, and a history of past reports, $h^t \equiv \sigma_t(\alpha, \theta^1, \theta^2, \dots, \theta^t, \tilde{h}^0, \dots, \tilde{h}^{t-1})$, where $\sigma \equiv \{\sigma_t(\cdot)\}$ is the player's strategy. With $h^0 \equiv \sigma_0(\alpha)$, $h^1 = \sigma_1(\alpha, \theta^1, \sigma_0(\alpha))$, \dots , h^t is determined recursively by σ on the path of play. To simplify notation, Θ_i^j and H_i^j will represent the ordered tuples $(\theta^i, \theta^{i+1}, \dots, \theta^j)$ and $(h^i, h^{i+1}, \dots, h^j)$ or, in a slight abuse of notation, the sets of such tuples.

The buyer utility type θ is private information in every round, and he is also privately informed about his type α before the contract is signed. There is a common prior over the distribution of the buyer types α . The seller has the power to commit to a contract and is risk neutral. Both buyer and seller are expected utility maximizers, and not liquidity constrained. Based on the revelation principle, the following derivations restrict attention to direct truthful mechanisms. The game starts in period 0, with the

² With infinite periods, a sufficient condition is that the overall payoff is the discounted sum of a per-period utility function that is uniformly bounded. That is, $\exists \delta \in [0, 1)$, $\exists B$ s.t. $\max_{\theta^t, q_t, p_t} |\delta^{-t} U_t(\theta^t, q_t, p_t)| < B$, which implies $\sum_t U_t(\cdot) < \infty$ (Fudenberg and Tirole 1991, p. 110).

buyer reporting his α value—in other words, his buyer type, after which he may have to make a payment, although we may assume *w.l.o.g.* that the payment is deferred to the subsequent rounds. In subsequent rounds, the buyer reports his θ utility type in each round, receives an allocation q and makes a new payment. In the following, the focus is on a game with infinite periods, but sometimes I may also consider a finite period game for description or as a stepping stone. The results will hold for the finite as well as infinite settings. In the finite case I will ignore discounting.

2.1 Reducing the number of ICC conditions

It is interesting to consider both the finite and the infinite time periods versions. For the infinite setting, a recursive approach could be attempted, 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 incentive compatibility 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 reduce this number as much as possible before proceeding to any analysis. The first step is proving that the relevant incentive compatibility conditions can be restricted. 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. With the initial period utility and the outside option at 0,

$$\forall \alpha : 0 \leq \sum_{t=1}^n \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(\alpha, \Theta_1^t), p_t(\alpha, \Theta_1^t)),$$

where $\Theta_1^t = (\theta^1, \dots, \theta^t)$ is one possible sub-history of utility types for period 1 to 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 . In addition, the buyer may consider participating in each subsequent round, generating other IRCs at each period $k \geq 1$:

$$\forall k \geq 1, \alpha, H_0^{k-1}, \theta^k : 0 \leq U_k(\theta^k, q_k(H_0^{k-1}, \theta^k), p_k(H_0^{k-1}, \theta^k)) + \sum_{t=k+1}^n \sum_{\Theta_k^t | \theta^k} \prod_{i=k+1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(H_0^{k-1}, \Theta_k^t), p_t(H_0^{k-1}, \Theta_k^t)),$$

where $H_0^{k-1} = (h^0, h^1, \dots, h^{k-1})$ is a history of past reports and $\Theta_k^t | \theta^k = (\theta^k, \dots, \theta^t)$ any following sub-history of true reports, conditional on the first report being θ^k . Because the seller has the flexibility to set the payment schedule, these additional IRCs will not constrain the result, implying that the buyer is not required

to commit for all time periods.³ The incentive compatibility constraints (ICCs) for a direct truthful mechanism are:

$$\begin{aligned} \forall \sigma, \alpha : & \sum_{t=1}^n \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(\alpha, \Theta_1^t), p_t(\alpha, \Theta_1^t)) \\ & \geq \sum_{t=1}^n \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(H_0^t), p_t(H_0^t)). \end{aligned} \tag{1}$$

where $H_0^n = (h^0, h^1, \dots, h^n)$ is given by σ . In particular, these expressions include the compatibility conditions for deviations at specific sub-histories only.

Proposition 2.1 *The set of incentive compatibility conditions (1) is satisfied if a stronger set of compatibility conditions (S) hold, whereby in each period and after any history 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, for $k \in \overline{0, n}$, $\forall \sigma, \alpha, \Theta_1^k$:*

$$\begin{aligned} & U_k(\theta^k, q_k(H_0^{k-1}, \theta^k), p_k(H_0^{k-1}, \theta^k)) \\ & + \sum_{t=k+1}^n \sum_{\Theta_k^t | \theta^k} \prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(H_0^{k-1}, \Theta_k^t), p_t(H_0^{k-1}, \Theta_k^t)) \\ & \geq U_k(\theta^k, q_k(H_0^k), p_k(H_0^k)) \\ & + \sum_{t=k+1}^n \sum_{\Theta_{k+1}^t} \prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(H_0^k, \Theta_{k+1}^t), p_t(H_0^k, \Theta_{k+1}^t)), \end{aligned}$$

where H_0^n is given by σ .

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. $\forall \sigma, \alpha, \Theta_1^n$:

$$U_n(\theta^n, q_n(H_0^{n-1}, \theta^n), p_n(H_0^{n-1}, \theta^n)) \geq U_n(\theta^n, q_n(H_0^n), p_n(H_0^n)). \tag{2}$$

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

³ See Proposition 3.1 for an example.

$$\begin{aligned} & \sum_{t=1}^n \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(H_0^t), p_t(H_0^t)) \\ & \leq \sum_{t=1}^{n-1} \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(H_0^t), p_t(H_0^t)) \\ & \quad + \sum_{\Theta_1^n} \prod_{i=1}^n \lambda_\alpha(\theta^i) U_n(\theta^n, q_n(H_0^{n-1}, \theta^n), p_n(H_0^{n-1}, \theta^n)), \end{aligned}$$

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. $\forall \sigma, \alpha$:

$$\begin{aligned} & \sum_{t=1}^n \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(H_0^t), p_t(H_0^t)) \\ & \sum_{t=1}^k \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(H_0^t), p_t(H_0^t)) \\ & \quad + \sum_{t=k+1}^n \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(H_0^k, \Theta_{k+1}^t), p_t(H_0^k, \Theta_{k+1}^t)). \end{aligned}$$

Assume that \mathcal{S} hold. \mathcal{S} say that, given 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.

$\forall \sigma, \alpha$:

$$\begin{aligned} & \sum_{t=1}^n \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(H_0^t), p_t(H_0^t)) \\ & \leq \sum_{t=1}^{k-1} \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(H_0^t), p_t(H_0^t)) \\ & \quad + \sum_{\Theta_1^k} \prod_{i=1}^k \lambda_\alpha(\theta^i) U_k(\theta^k, q_k(H_0^k), p_k(H_0^k)) \\ & \quad + \sum_{t=k+1}^n \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(H_0^k, \Theta_{k+1}^t), p_t(H_0^k, \Theta_{k+1}^t)) \end{aligned}$$

$$\begin{aligned}
 &= \sum_{t=1}^{k-1} \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(H_0^t), p_t(H_0^t)) \\
 &\quad + \sum_{\Theta_1^k} \prod_{i=1}^k \lambda_\alpha(\theta^i) \left[U_k(\theta^k, q_k(H_0^k), p_k(H_0^k)) \right. \\
 &\quad \left. + \sum_{t=k+1}^n \sum_{\Theta_1^t | \Theta_1^k} \prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(H_0^k, \Theta_{k+1}^t), p_t(H_0^k, \Theta_{k+1}^t)) \right] \\
 &\leq \sum_{t=1}^{k-1} \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(H_0^t), p_t(H_0^t)) \\
 &\quad + \sum_{\Theta_1^k} \prod_{i=1}^k \lambda_\alpha(\theta^i) \left[U_k(\theta^k, q_k(H_0^{k-1}, \theta^k), p_k(H_0^{k-1}, \theta^k)) \right. \\
 &\quad \left. + \sum_{t=k+1}^n \sum_{\Theta_1^t | \Theta_1^k} \prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(H_0^{k-1}, \Theta_k^t), p_t(H_0^{k-1}, \Theta_k^t)) \right] \\
 &= \sum_{t=1}^{k-1} \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(H_0^t), p_t(H_0^t)) \\
 &\quad + \sum_{t=k}^n \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(H_0^{k-1}, \Theta_k^t), p_t(H_0^{k-1}, \Theta_k^t)).
 \end{aligned}$$

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

$$\sum_{t=k}^{\infty} \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(H_0^t), p_t(H_0^t)), \tag{3}$$

which is bounded and must go to 0 as $k \rightarrow \infty$, implicit if the infinite sums in the ICC conditions are well defined.⁴ For an infinite direct mechanism that implements a menu, say that the infinite set of constraints \mathcal{S} will hold. Let σ represent an optimal but not necessarily truthful strategy, mapping any buyer type and sequence of utility types to optimal reports, and let τ be the truthful strategy. The proof idea is to devise an approximate strategy that is a “truncated” version of σ , in which the buyer reports

⁴ For example, it is sufficient to express U_t as a weakly concave utility function multiplied with geometric discounting, and assume an upper bound on the money transfers in any period.

according with the strategy σ up to period k , and then he reports truthfully, so according to τ . Let σ_k be that strategy. 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, i.e., $V_{\sigma_k} \rightarrow V_\sigma$ as $k \rightarrow \infty$.⁵ 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, i.e., $\forall k \geq 0: V_\tau \geq V_{\sigma_k}$. But by assumption $V_\tau \leq V_\sigma$, so letting k go to infinity gives us by dominated convergence that $V_\tau = V_\sigma$. We obtain that truthful reporting is weakly optimal in the context of a menu satisfying \mathcal{S} , which implies that that \mathcal{S} are more restrictive than the ICCs, so the regular ICCs have to hold when \mathcal{S} hold. \square

Intuitively, the alternative conditions of Proposition 2.1 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. Once there, it follows that \mathcal{S} restrict the possible mechanisms at least as much as the ICCs, since the truthful strategy is optimal only if the ICCs hold.

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 must be true as long as his lifetime utility expressed as an infinite sum over periods is well defined; then the partial sum over time 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 2.2 *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 he is in periods after the 0th, considers only one-time deviations from truth, followed by truthful reporting in the future. That is, the ICCs imply the following inequalities, for⁶ $k \in \overline{0, n}$, $\forall \sigma, \alpha, \Theta_1^k$:*

⁵ V_σ is a shorthand for the value of expression (3) evaluated for strategy σ .

⁶ At $k = 0$, $U_0 = 0$ as defined.

$$\begin{aligned}
 &U_k(\theta^k, q_k(\alpha, H_1^{k-1}, \theta^k), p_k(\alpha, H_1^{k-1}, \theta^k)) \\
 &+ \sum_{t=k+1}^n \sum_{\Theta_1^t | \Theta_1^k} \prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(\alpha, H_1^{k-1}, \Theta_k^t), p_t(\alpha, H_1^{k-1}, \Theta_k^t)) \\
 &\geq U_k(\theta^k, q_k(\alpha, H_1^k), p_k(\alpha, H_1^k)) \\
 &+ \sum_{t=k+1}^n \sum_{\Theta_1^t | \Theta_1^k} \prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(\alpha, H_1^k, \Theta_{k+1}^t), p_t(\alpha, H_1^k, \Theta_{k+1}^t)),
 \end{aligned}$$

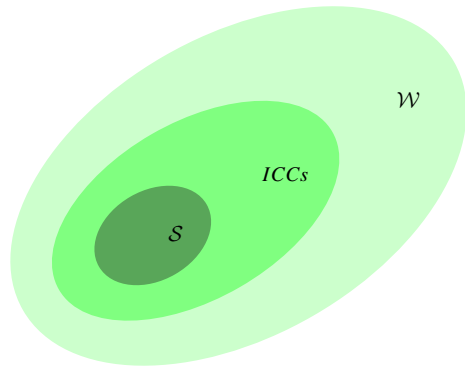
and

$$\begin{aligned}
 &\sum_{t=1}^n \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(\alpha, \Theta_1^t), p_t(\alpha, \Theta_1^t)) \\
 &\geq \sum_{t=1}^n \sum_{\Theta_1^t} \prod_{i=1}^t \lambda_\alpha(\theta^i) U_t(\theta^t, q_t(h^0, \Theta_1^t), p_t(h^0, \Theta_1^t)).
 \end{aligned}$$

Proof The proof is obvious, and 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 (α, H_1^{k-1}) , when $\sigma_k(\alpha, \Theta_1^{k-1}, \theta^k) = h^k$ and, subsequently, again truthful reporting after any continuation history. One can simplify the ICC by eliminating the payoffs for the histories $\Theta_1^n - \Theta_1^n | \{(\alpha, H_1^{k-1}, \theta^k)\}$ and their continuations, which must be equal, to get:

$$\begin{aligned}
 &\prod_{i=1}^k \lambda_\alpha(\theta^i) U_k(\theta^k, q_k(\alpha, H_1^{k-1}, \theta^k), p_k(\alpha, H_1^{k-1}, \theta^k)) \\
 &+ \prod_{i=1}^k \lambda_\alpha(\theta^i) \sum_{t=k+1}^n \sum_{\Theta_1^t | (\alpha, H_1^{k-1}, \theta^k)} \\
 &\prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(\alpha, H_1^{k-1}, \Theta_k^t), p_t(\alpha, H_1^{k-1}, \Theta_k^t)) \\
 &\geq \prod_{i=1}^k \lambda_\alpha(\theta^i) U_k(\theta^k, q_k(\alpha, H_1^k), p_k(\alpha, H_1^k)) \\
 &+ \prod_{i=1}^k \lambda_\alpha(\theta^i) \sum_{t=k+1}^n \sum_{\Theta_1^t | (\alpha, H_1^{k-1}, \theta^k)} \\
 &\prod_{j=k+1}^t \lambda_\alpha(\theta^j) U_t(\theta^t, q_t(\alpha, H_1^k, \Theta_{k+1}^t), p_t(\alpha, H_1^k, \Theta_{k+1}^t))
 \end{aligned}$$

Fig. 1 The restricted domains according to \mathcal{W} , \mathcal{S} , $ICCs$ in the menu parameter space



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 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 $\mathcal{W} \subset ICCs$, and $\mathcal{W} \subset \mathcal{S}$ is by definition. Proposition 2.1 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} , 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 $\mathcal{D}[\mathcal{S}] \subset \mathcal{D}[ICCs] \subset \mathcal{D}[\mathcal{W}]$ (Fig. 1). 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 $\mathcal{D}[ICCs]$. A containing domain may generate a false maximum, if it is located outside $\mathcal{D}[ICCs]$, while a subdomain can lead to maxima on its boundaries that could be suboptimal. To solve for the optimal mechanism in the application that follows, my strategy will be to find the optimal menu in $\mathcal{D}[\mathcal{W}]$ and show that it is in $\mathcal{D}[\mathcal{S}]$.

The private information of the buyer is separated into two parts: the type α , which is private knowledge ex-ante, and the type θ , which is conditionally distributed according to α and known after the contract is signed. Intuitively, this methods relies on the fact that we can always segregate the private information of the buyer in such a way. After that, \mathcal{W} is the set of constraints that insure truthful reporting once α is known, so any mechanism constrained by \mathcal{W} will be optimal as long as the contract has payments at time 0 which incentivize the truthful reporting of α . Since the value of the relationship

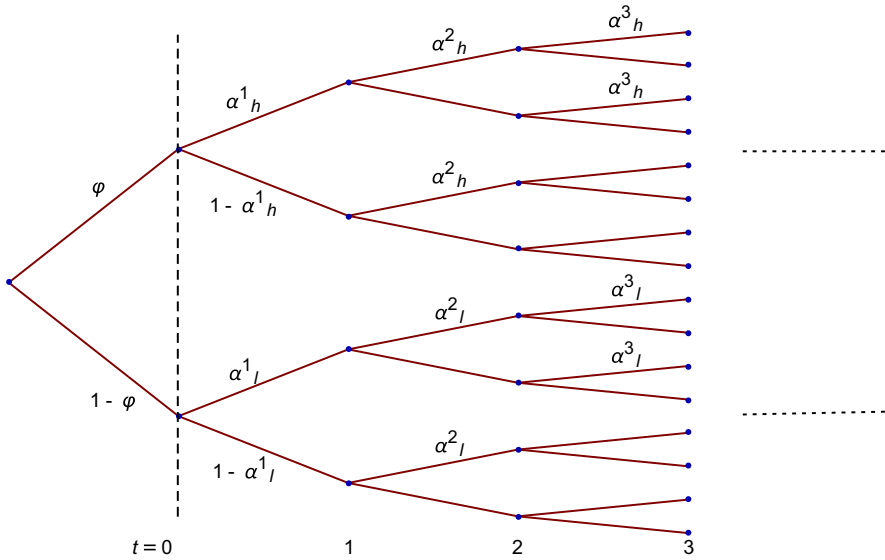


Fig. 2 Moves by nature determining types. The vertical dotted line shows when the contract is signed

between the buyer and the seller conditional on α is common knowledge, this is always possible.

3 An application with two types

In the following application, the buyer is of two possible types, and has two possible utility types. The buyer’s utility in any one time period is given by

$$U_t(\theta^t, q_t, p_t) = \theta^t q_t - (q_t)^2/2 - p_t(q_t),$$

where q_t is the buyer’s consumption in period t , p_t is the price he pays for it, henceforth determined by q_t . θ^t is the utility or demand type, and can take one of the values $\{\theta_h^t, \theta_l^t\}$, which are common knowledge.⁷ The buyer has ex-ante private information about the probability of having a utility type θ_h^t instead of θ_l^t in any one period. This probability is determined by the buyer’s type, and can take the values $\{\alpha_h^t, 1 - \alpha_h^t\}$ or $\{\alpha_l^t, 1 - \alpha_l^t\}$, also common knowledge. Using the notation of the general set-up, $\alpha_h^t \equiv \lambda_h(\theta_h^t)$, $1 - \alpha_h^t \equiv \lambda_h(\theta_l^t)$, $\alpha_l^t \equiv \lambda_l(\theta_h^t)$, and $1 - \alpha_l^t \equiv \lambda_l(\theta_l^t)$. The common prior is that the high and low buyer types appear with odds $\{\varphi, 1 - \varphi\}$ respectively. This setting is sketched in Fig. 2. The game starts in period 0, with the buyer reporting his buyer type—so future α values, and he reports his utility type in each subsequent round—the θ value, receiving his allocation q^t and making a new payment. For all t ,

⁷ It is common to consider the quadratic portion, $q^2/2$, as a cost to the seller for producing quantity q . If we define the price net of cost, it does not change the analysis or interpretation.

let $\theta_h^t > \theta_l^t \geq 0$, $1 > \alpha_h^t > \alpha_l^t > 0$, $1 > \varphi > 0$.⁸ To relate to literature, observe that $\alpha_h^t > \alpha_l^t$ implies that the draw from $\{\alpha_h^t, 1 - \alpha_h^t\}$ first-order stochastically dominates the draw from $\{\alpha_l^t, 1 - \alpha_l^t\}$. The initial results in Sect. 2.1 were derived without making assumptions on $U_t(\theta^t, q_t, p_t)$ and only requiring that all θ^t are from the same discrete set. In the following, the time index will be dropped for clarity. The arguments are the same for all periods considered, but the values of θ and α are allowed to change over time provided that the inequalities and the common knowledge assumptions above hold.

3.1 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. The goal is to obtain specific expressions for allocations, and to interpret and describe the equilibrium. This section shows through a few lemmas that the optimal menu must have first best allocations for reports of θ_h , as well as for θ_l after type α_h was reported. The following lemmas hold 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 one 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

⁸ Strict inequalities are assumed to consider interesting cases. The results hold in the limit.

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:

$$\begin{aligned} &\theta_h(q_h + \epsilon) - (q_h + \epsilon)^2/2 - (p_h + (\theta_h - q_h)\epsilon - \epsilon^2/2) + u_h \\ &= \theta_h q_h - q_h^2/2 - p_h + u_h. \end{aligned}$$

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

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

\Leftrightarrow

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

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:

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

⇔

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

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. □

This lemma shows that the θ_h type must always be allocated at least the first best in an optimal mechanism, because one 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. □

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. □

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, one 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.

3.2 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 an α_h buyer pretending to have the α_l type must be subtracted with the appropriate weight.

Lemma 3.5 *Let v be the α_l -type buyer value at the beginning of any period $t \geq 1$, v_h and v_l the buyer continuation value after a θ_h^t and θ_l^t report respectively, and $P(v)$, $P(v_h)$ and $P(v_l)$ ⁹ the corresponding seller values. The simplified seller’s maximization problem can be recursively stated as:*

$$\begin{aligned}
 P(v) = & \max_{\{p_h, p_l, v_h, v_l, q_l\}} (1 - \varphi)[\alpha_l p_h + (1 - \alpha_l) p_l] + \alpha'_l \delta P(v_h) + (1 - \alpha'_l) \delta P(v_l) \\
 & + \varphi[-\alpha_h(\theta_h^2/2 - p_h) - (1 - \alpha_h)(\theta_l q_l - q_l^2/2 - p_l)] + C \tag{4}
 \end{aligned}$$

s.t. (i) $\theta_h^2/2 - p_h + \delta v_h = \theta_h q_l - q_l^2/2 - p_l + \delta v_l$,
(ii) $v = \alpha_l(\theta_h^2/2 - p_h) + (1 - \alpha_l)(\theta_l q_l - q_l^2/2 - p_l) + \alpha'_l \delta v_h + (1 - \alpha'_l) \delta v_l$,
(iii) $q_l \geq 0$.

Proof Since the behavior of the α_h type is already determined, the transfers and allocations on the α_h branch are of no concern, except for the fact that the seller pays information rent to the α_h type based on the menu offered on the α_l branch. Specifically, with probability φ , the seller loses $\alpha_h(\theta_h^2/2 - p_h) + (1 - \alpha_h)(\theta_l q_l - q_l^2/2 - p_l)$ in expectation to the α_h type as information rent. With probability $1 - \varphi$, he gains the expected payment of the α_l type, $\alpha_l p_h + (1 - \alpha_l) p_l$.¹⁰ C minus the period specific information rent is the seller’s gain from the α_h type. Hence, the objective function is the expected sum in (4).

By assumptions and \mathcal{W} , the buyer reports θ truthfully if the constraint (i) holds. Before having the period-specific private information, he current value of the buyer is the expected gain from the trade, $\alpha_l(\theta_h^2/2 - p_h) + (1 - \alpha_l)(\theta_l q_l - q_l^2/2 - p_l)$, plus the discounted future expected value, $\alpha'_l \delta v_h + (1 - \alpha'_l) \delta v_l$, so (ii) holds. The allocations can only be positive, so (iii) holds. □

The optimal mechanism chooses p_h, p_l, v_l, q_l to maximize seller value (including the continuation value). The discounted expected continuation value is $\alpha'_l \delta P(v_h) + (1 - \alpha'_l) \delta P(v_l)$ (Fig. 3). Since we assumed that the α_l -type rationality constraint in period

⁹ Time indices are dropped for simplicity. Future type probabilities are denoted by a prime symbol.

¹⁰ While the seller is now informed of the α type, he has to commit to the contract *ex ante*.




Fig. 3 Node sets determining buyer values in the recursive formulation, for the first period with allocations

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 3.1 *Assuming that condition (iii) in Lemma 3.5 does not bind, the allocation of a low buyer type after reporting θ_l in the maximization problem (4) 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 $v = 0$ in the first period and with the normalization $v_h = v_l = 0$ in all periods,

$$p_h = \frac{\theta_h^2}{2} - (\theta_h - \theta_l)(1 - \alpha_l)q_l,$$

$$p_l = (\alpha_l(\theta_h - \theta_l) + \theta_l)q_l - \frac{q_l^2}{2}.$$

(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 that period.

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

$$p_h = -v + \delta v_h + \frac{\theta_h^2}{2} - q_l(1 - \alpha_l)(\theta_h - \theta_l),$$

$$p_l = -\frac{q_l^2}{2} - v + v_l\delta + q_l(\alpha_l(\theta_h - \theta_l) + \theta_l).$$

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, so changing 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. □

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

Proposition 3.2 *Assuming that condition (iii) in Lemma 3.5 does not bind, 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. When (iii) in Lemma 3.5 holds with equality in a period, there is no more transfer from the seller, and the high buyer type allocations and prices in that period are,

$$q'_h = \frac{\theta_h^2}{2}, \quad q'_l = \theta_l - (\theta_h - \theta_l) \frac{\alpha_h}{1 - \alpha_h},$$

$$p'_h = \frac{q_h^2}{2} + (\theta_h - \theta_l)q_l, \quad p'_l = \frac{q_l^2}{2}.$$

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 straightforward 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.

When the non-negativity constraint (iii) binds, there is no more occasion for the α_h type to earn information rents from misrepresenting his type, because $q_l = 0$ and the α_l type pays for all the value, whichever θ he reports. At this point, the seller should find it optimal to start distorting the θ_l allocation of the α_h buyer type. Formally, because the compatibility condition for the θ_h type binds, Lemma 3.4 does not imply $q'_l \geq \theta_l$ anymore. One can solve the problem as if only the α_h type exists in the standard way, and check the compatibility conditions. It is easy to verify that the seller will extract the entire value, and with 0 transfers at time 0, and the allocations from Proposition 3.1, all the incentive compatibility constraints are satisfied. \square

Observe that the allocations and prices depend on the period only through the values of the parameters $\{\alpha_{h,l}^t, \theta_{h,l}^t\}$, which justifies dropping the time indices. It is left to verify 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 to check are the conditions in period 0. The IRCs are satisfied by construction. If an α_h -type buyer reports α_l^1 , then the information rent he earns is given by:

$$v_h = (\alpha_h^1 - \alpha_l^1)(\theta_h^1 - \theta_l^1)q_l + \delta\alpha_h^2v_h + \delta(1 - \alpha_h^2)v_h,$$

which solves to

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

This is the transfer an 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 :

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

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

Reintroducing the time indices for parameters that were dropped in the derivations, in Proposition 3.1:

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

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

Varying the different parameters gives us results that coincide with those of simpler problems. If $\varphi = 0$, then there is only one buyer type, who must sign a contract before he knows his first θ . As expected, our solution becomes $q_l^t = \theta_l^t$, $q_h^t = \theta_h^t$, which is the classical dynamic contract when the buyer has no private information. Substituting in the prices determined in Proposition 3.1, it is readily seen that

$$\alpha_l^t p_h^t + (1 - \alpha_l^t) p_l^t = \alpha_l^t \frac{\theta_h^{t2}}{2} + (1 - \alpha_l^t) \frac{\theta_l^{t2}}{2}, \quad (5)$$

$$p_h^t - p_l^t = \frac{\theta_h^{t2}}{2} - \left(\theta_h^t \theta_l^t - \frac{\theta_l^{t2}}{2} \right). \quad (6)$$

With no private information, the buyer is allocated the first best, and pays an expected price equal to the expected value of the allocation. The difference in the prices for the two allocations reflects the information rent—the potential gain from lying about the utility type.

As φ approaches 1, q_l^t becomes 0 when condition (iii) in Lemma 3.5 binds. From Proposition 3.2, the low buyer type is allocated $q_l^t = 0$ after low utility reports, and the high buyer type receives reduced allocations q_h^t after low reports. The rent becomes 0 for the high buyer type because he cannot gain anymore from misrepresenting α . In effect, the high type buyer is so dominant that the other type is distorted to the limit, and the high type is treated as in the classic problem with only one type and no private information.

When $\alpha_h^t = \alpha_l^t$, in the limit the solution is the first best because the buyer has no relevant private information that is not known at the moment of contracting. In this case, the buyer types are not distinctive anymore, a situation equivalent to equivalent to $\varphi = 0$.

The dynamics in this multi-period model are the result of the time variation of the parameters $\{\theta_{l,h}^t\}$ and $\{\alpha_{l,h}^t\}$, which are common knowledge. Since both the two utility types and their distribution are allowed to change over time, allocations can vary as well. However, the same expression derives the allocations in each period after the first. Keeping in mind that the allocation distortions have the purpose to elicit truthful reporting of private information, such private information is revealed through the buyer type initially and then through the utility type every subsequent period. The expression for the allocation distortions is common because the same type of private information is revealed. In other words, there is no reason to distort n th period allocations differently from $(n + 1)$ th period allocations—everything else being equal. This reasoning holds only if the true buyer type is elicited at onset; otherwise each period's utility type report would lead to an updated prior for the buyer type as well.

It is instructive to contrast the dynamic evolution of allocations with that in Battaglini (2005). The difference is in the informativeness of the private information the agent has when the contract is signed. When the type follows a Markov process, the initial type report has an exponentially decreasing correlation to the future type; hence, the distortions decrease over time and the contract converges to the contract that would be offered to a buyer with no private information (i.e., no distortions, and rents are extracted *ex ante*.) In the current setting, the initial buyer type report is relevant for

all subsequent type reports, so the distortions are consistent over time. One can see the current setting as repeating a one period game infinitely many times, with utilities and payments discounted. Since the only connection between periods comes from the buyer type, which is reported in the truthful mechanism, the optimal allocations in all future periods—so also the distortions from the first best—have to be similarly derived. When types are given by a Markov process, the slow reduction in future allocation distortions takes into account the progressively weaker impact of the initial private information, which is the result of the geometrically decreasing correlation between types over time. The form of the allocation distortions at all time periods is induced by the trade-off between information rent and profit, so it is related to the impact of the private information that the buyer or agent has at the moment when the contract is signed.¹¹

In future work, a straightforward but involved extension is solving for an arbitrary finite number of buyer types. The equivalents of Lemmas 3.1—3.4 can be stated for every pair of adjacent types, and then Lemma 3.5 must be reformulated for a marginally distorted type—equivalent in our case to the α_l type. To simplify and reduce compatibility constraints further, a new condition on the distribution of types is required to make sure that only incentive constraints between adjacent types are relevant. This condition would become, in the limit of continuous types, proportional to the inverse hazard rate. A conjecture is that for $\{\alpha_i\}_{i=1}^I$ there are two indices $1 \leq x < y \leq I$ s.t. for the α_j buyer type with $x \leq j \leq y$, the low utility allocation is reduced, but not entirely eliminated, and the reduction is now proportional to the probability fraction of all the types $i > j$ over the probability of type α_j , so $\sum_{j+1}^I \lambda_{\alpha_i} / \lambda_{\alpha_j}$. For α_i , $i < x$, the low utility allocation is 0, while for α_i , $i > y$, the low utility allocation is efficient. This also hints at the solution in the continuous limit, where α_x , α_y would have to be the bounds of a subinterval of types where low utility allocations are distorted. Lower types would then receive 0 at low utility, and higher types receive the first best. The continuous limit of the ratio of probabilities is then the inverse hazard rate, divided by a scale factor.¹²

5 Conclusion

This paper analyses finite and infinite time contracts between a principal and an agent, specifically focusing on set-ups in which the agent types and their utility types are discrete, and when the agent has private information on the probability of his future utility types. In general, the number of compatibility conditions that need to be considered is very large, but the solution can be searched in a wider space bounded by fewer or weaker constraints, and confirmed with a stronger set of constraints. The paper proposes a way to construct these weaker and stronger constraints. First look for direct truthful mechanisms that incentivize a buyer who has reported his type truthfully to

¹¹ If the buyer has no relevant private information about a future period at the moment of signing the contract, the principal can always ask for the expected information rent as a fee.

¹² A scale factor must appear since, if the range of the type distribution is “stretched”, the probability density and hazard rate are reduced. However, the distortion should stay the same, since the buyer type doesn’t directly influence utility.

also report his utility type truthfully, after any history of reports. Then, verify that the mechanism implies that all types will respond truthfully, regardless of history.

The method can be applied when the agent has private information that is not completely random nor fixed, but comes from a type dependent distribution which is his private information. In an example with two discrete types each period, the optimal contract for two types of buyers with two types of utilities has some of the usual properties found in the static problem, like no allocation distortions for the high type. There are consistent distortions for the low type buyer when he reports low utility, and the distortions decrease social welfare. The permanence of distortions reflects the persistent value of the initial private information. The results provide some insight into the pricing strategies of monopolistic suppliers of firms, franchises, or sellers that have repeated interactions with their customers.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Athey S, Segal I (2013) An efficient dynamic mechanism. *Econometrica* 81(6):2463–2485
- Baron DP, Besanko D (1984) Regulation and information in a continuing relationship. *Inf Econ Policy* 1(3):267–302
- Battaglini M (2005) Long-term contracting with markovian consumers. *Am Econ Rev* 95:637–658
- Battaglini M (2007) Optimality and renegotiation in dynamic contracting. *Games Econ Behav* 60(2):213–246
- Bergemann D, Välimäki J (2010) The dynamic pivot mechanism. *Econometrica* 78(2):771–789
- Boleslavsky R, Said M (2013) Progressive screening: long-term contracting with a privately known stochastic process. *Rev Econ Stud* 80(1):1–34
- Courty P, Li H (2000) Sequential screening. *Rev Econ Stud* 67(4):697–717
- Cremer J, McLean RP (1988) Full extraction of the surplus in bayesian and dominant strategy auctions. *Econ J Econ Soc*, pp 1247–1257
- Dewatripont M (1989) Renegotiation and information revelation over time: the case of optimal labor contracts. *Q J Econ*, pp 589–619
- Eső P, Szentes B (2007) Optimal information disclosure in auctions and the handicap auction. *Rev Econ Stud* 74(3):705–731
- Fudenberg D, Tirole J (1991) *Game theory*. MIT Press. ISBN 9780262061414
- Hart OD, Tirole J (1988) Contract renegotiation and coasian dynamics. *Rev Econ Stud* 55:509–540
- Laffont JJ, Tirole J (1988) The dynamics of incentive contracts. *Econo J Econ Soc*, pp 1153–1175
- Laffont JJ, Tirole J (1990) Adverse selection and renegotiation in procurement. *Rev Econ Stud* 57(4):597–625
- Laffont JJ, Tirole J (1996) Pollution permits and compliance strategies. *J Public Econ* 62(1):85–125
- Miravete EJ (2003) Choosing the wrong calling plan? ignorance and learning. *Am Econ Rev* 93:297–310
- Pavan A, Segal I, Toikka J (2014) Dynamic mechanism design: a myersonian approach. *Econometrica* 82(2):601–653
- Rustichini A, Wolinsky A (1995) Learning about variable demand in the long run. *J Econ Dyn Control* 19(5):1283–1292