ESSAYS IN DYNAMIC MECHANISM DESIGN

Rohit Lamba

A DISSERTATION

PRESENTED TO THE FACULTY

OF PRINCETON UNIVERSITY

IN CANDIDACY FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE

BY THE DEPARTMENT OF ECONOMICS

ADVISOR: Stephen Morris

JUNE 2014

To my parents: Kapil and Seema Lamba

Abstract

Questions of design in real economic situations are often dynamic. Managerial compensation, repeated auctions, and taxation are good examples. These demand the economic theory of mechanism design to be adept to changing underlying environments and evolving information. Adjusting existing static results to the dynamic models and introducing new ones is thus what the doctor orders. This collection of essays is a contribution to the theory and applications of dynamic mechanism design.

Chapter 1 asks the question: *when can efficient institutions be made self enforcing?* To answer it, the setting of bargaining with two sided asymmetric information is chosen– a buyer has a hidden valuation for a good and a seller can produce the good at a hidden cost, both of which can change over time. The essay provides necessary and sufficient conditions for efficiency in this bilateral trading problem. In the process of establishing this result, a new notion of budget balance is introduced that allows the budget to be balanced dynamically, borrowing from the future but in a bounded fashion. Through a set of simple examples the comparative statics of the underlying economics forces of discounting and level of asymmetric information are explored.

In chapter 2, a *dynamic and history dependent version of the payoff equivalence* result is established. It provides an equivalence class of all mechanisms that are incentive compatible. Given two mechanisms that implement the same allocation, expected utility of an agent after any history in one must differ from the other through a history dependent constant. This result is then exploited to unify a host of existing results in efficient dynamic mechanism design. In particular a mechanism, and necessary and sufficient conditions are provided for the implementation of the efficient allocation in a general N-player dynamic mechanism design problem under participation constraints and budget balance.

Finally, in chapter 3 (coauthored with Marco Battaglini), we explore the *applicability and limitations of the first-order approach in solving dynamic contracting models,* *and the nature of contracts when local constraints are not sufficient to characterize the optimum.* A dynamic principal-agent model in which the agent's types are serially correlated forms the backbone of the analysis. It is shown that the first-order approach is violated in general environments; when the time horizon is long enough and serial correlation is sufficiently high, global incentive compatibility constraints generically bind. By fully characterizing a simple two period example, we uncover a number of interesting features of the optimal contract that cannot be observed in special environments in which the standard approach works. Finally, we show that even in complex environments, approximately optimal allocations can be easily characterized by focusing on a class of contracts in which the allocation is forced to be monotonic.

Acknowledgements

One of my favorite poets, Rumi, once wrote[1]

Mind does its fine-tuning hair-splitting
but no craft or art begins
or can continue without a master
giving wisdom into it

It is with the same emotion I express a deep gratitude to my teacher, mentor and advisor, Stephen Morris. His dedication towards the craft of economic theory and his students is priceless. It was his course on mechanism design that got me hooked on to the subject, and it is the long hours spent in his office that have been pivotal in creation of this dissertation that you are about to read.

Marco Battaglini has taught me the foundations of dynamic mechanism design. At the beginning of the research stage, he nudged me towards understanding the still nascent literature, culminating into our joint project that forms the third essay. This dissertation would not exist without him. Sylvain Chassang has an infectious energy for research that he effortlessly passes on to budding economists. He provided critical feedback and constant encouragement at crucial moments of my research. Dilip Abreu has been a rock solid mentor since my early days at Princeton. It would not be an exaggeration to say that his legend inspired me to do a PhD in economic theory. The inspiration has only grown as I've got to know him personally. Wolfgang Pesendorfer and Faruk Gul were generous with their time, feedback and encouragement. I'm immensely grateful to all of these amazing professors.

I took six months off during my PhD. During this time I had the privilege of working with Raghuram Rajan at the Office of the Chief Economic Advisor in Delhi.

[1]The Soul of Rumi, translations by Coleman Banks, Harper Collins 2001

His mentorship and enthusiasm have had an indelible impact on my approach to the academic thought process. Many thanks go out to him.

My first year study group- Jacob Goldin, Nikolaj Harmon, Tatiana Homonoff and Peter Buisseret- was critically important. We were a constant source of emotional and intellectual support to each other. I owe them a huge debt. Jacob Goldin especially has been a treasure of a friend. His generosity, kindness and intellectual spirit made my time at Princeton memorable and enjoyable. I'd also like to thank Nemanja Antic, Ben Brooks, Karyne Charbonneau, Gonzalo Cisternas, Diogo Guillen, Josephine Duh, Jay Lu, Debi Schwartz and Phyllis Sun for their encouragement and friendships. Karyne provided love and support that I'll always be grateful for.

Rohit De, my namesake, was an instrumental part of my time at Princeton; a roommate par excellence, and a true friend in need and deed. Arijeet Pal grounded me, helped me grow and brought much happy nostalgia of being away from home. Vinay Sitapati taught me to keep my work always connected to the changing social landscapes of the real world. Wamiq Umaira provided critical feedback on the vagaries of life and was a constant support. Darren Pais and Dinsha Mistree were generous in their friendships. Azza Cohen has inspired me to dream, love, think and care. Her contribution to my life in Princeton is immeasurable. A big thank you to them all.

Lastly, I'd like to thank my family. My sister has always believed in me. Her unconditional love has always been my anchor. My Nanaji gave me the first lessons in mathematics, and introduced me to the wonders of logical thinking. He would have been proud to see this dissertation. My parents have taught me to dream big, take risks and follow my passions. They are the reason of my being in every sense. I dedicate this dissertation to them.

Contents

Chapter 1

Efficiency in Repeated Bargaining: A Mechanism Design Approach

1.1 Introduction

In a paper that would generate much interest amongst economists and legal scholars, Coase [1960] argued that if transactions costs are low enough and trade a possibility, bargaining will eventually lead to an efficient outcome independent of the initial distribution of property rights. A few decades later, in perhaps an equally influential paper, Myerson and Satterthwaite [1983] showed that under reasonable institutional assumptions, asymmetric information precludes efficient trade. A key missing link in Coase's argument was established as part of the growing acceptance of the role of information in economic transactions.[1]

[1]Introducing the *Myerson and Satterthwaite Theorem*, as now it is popularly called, Milgrom [2004] writes

> "Doubts about the [Coase's] efficiency axiom are based partly in concern about bargaining with incomplete information. After all, a seller is naturally inclined to exaggerate the cost of his good, and a buyer is inclined to pretend that her value is low. Should we not expect these exaggerations to lead sometimes to missed trading opportunities?"

Myerson and Satterthwaite [1983] , along with many other papers that came before and after, asked important questions of institution design under varying objectives- efficiency, revenue maximization, etc. Public goods provision, procurement, auctions, optimal taxation, bilateral trading, wage contracts are just some applications of the general theory of mechanism design that has thus developed.[2]

Most of these papers, including the two aforementioned, dealt with static or one-time interactions. Arguably, many of these economic transactions are inherently dynamic, where information revealed today can be used for contract design tomorrow. Food subsidies are provided repeatedly. In a fast changing technological landscape, spectrum auctions and buybacks are taking place repeatedly. Taxation is often dynamic and tagged with age, social security being a case in point. Wage contracts and bonuses depend on performance parameters evaluated over time. Online selling can now rely on a huge treasure trove of past buying data.

This paper seeks to provide a theory of such dynamic institutions and contribute towards the burgeoning literature on dynamic mechanism design.[3] When are efficient institutions self enforcing? There are three key words in the preceding statement. By, efficiency we mean first-best or the optimal allocation of resources without any additional frictions or binding constraints. Institution is an environment characterized by a set of rules that internalize underlying frictions. And, self enforcing, refers to the ability of the institution to implement the desired objective under limits on external subsidies.

To answer this question, we choose the well studied static problem of bargaining under two-sided asymmetric information that concerned Myerson and Satterthwaite [1983]. A seller wants to repeatedly sell a non-durable good to a buyer. Their valuations for the good are privately known and can change over time. Repetition

[2]See Mas-Collel, Whinston and Green [1995], and Jackson [2003] for a thorough overview of the literature.

[3]See Bergemann and Said [2010] and Vohra [2012] for insightful surveys.

can blunt the impossibility of efficiency result of Myerson and Satterthwaite [1983]. We formalize the extent and logic of this “blunting".

The main result of the paper is to provide a necessary and sufficient condition for the implementation of efficient trade under participation constraints and budget balance. The interaction of these two constraints with private information leads to the impossibility of efficiency result in the static framework. We model the repeated interaction of these constraints and precisely characterize when the impossibility result can be overturned.

The key conceptual contribution we bring to the table is that of *interim budget balance*.[4] The motivation is to allow for the role of a financial intermediary or mechanism designer, but one who cannot have access to an unbounded credit line. At any given history, the expected value of current and future cash flows from the buyer and the seller must be non-negative. Interim budget balance can be seen as the mechanism design counterpart to self enforcing constraints from the relational contracting literature[5], but not on the side of the agents, rather the institution itself; and standard bond issuing deficit financing constraints in macro models.[6]

In the process of providing a tight characterization of efficiency, we construct a dynamic and modified version of the Vickery-Clarkes-Groove mechanism which provides the mechanism designer the maximal surplus every period, and the minimum utility to the agents subject to their information rents and reservation values. By construction, if this mechanism does not produce an expected budget surplus at every history of the contract, and hence satisfy interim budget balance, no other mechanism does. Also, in a very simple implementation of this mechanism, each agent pays a small fees at the start of every period, post which the mechanism designer runs a

[4]The standard notions in the literature are that of ex ante budget balance- aggregate ex ante expected cash flow to the mechanism designer is non-negative, and ex post budget balance- transfers sum to zero every period for any history.

[5]See Thomas and Worrall [1988], and Levin [2003]

[6]See Ljungqvist and Sargent [2004].

(static) VCG mechanism.

We choose the two types example to illustrate some of the key economic forces at play. The likelihood of achieving efficiency is increasing function of the discount factor and a decreasing function of the level of asymmetric information in the model, as measured by persistence of Markov process. In fact, as both discounting converges to one and the Markov process converges to constant types, the attainment of the efficiency depends on the order of limits– can be attained if the former converges faster, but cannot be attained if the speed of convergence of the latter is higher.

A key advantage of looking at the bargaining model through a dynamic framework is that allows for the role of an intermediary who is forward looking and can help the agents reach an efficient outcome while possibly earning some rents as well. In a world with no credit constraints, the agents can perfectly insure each other against bad shocks as long as long types are not constant and the discount factor is high. However, for high levels of persistence the transfers required for this insurance can get large after bad shocks. This creates a role for an intermediary. We also show that the ability of the intermediary to induce the efficient implementation rises with access to a savings technology.

Related Literature. The paper builds on a body of work from static mechanism design. Vickery [1961], Groves [1973] and d'Aspremont and Gérard-Varet (1979) were some of the early papers to talk about efficient mechanisms. The bilateral trading problem we study has a rich tradition in the static mechanism design literature- Myerson and Satterthwaite [1983], and Chatterjee and Samuelson [1983] being two of the early papers. Myerson and Satterthwaite [1983 established the impossibility of reconciling the three key constraints of private information, participation constraints and budget balance with the efficient allocation. Williams [1999] and Krishna and Perry [2000] prove the same result by a different technique, exploiting the VCG mechanism.

Athey and Segal [2007b, 2013] generalize the AGV mechanism to the dynamic model. The mechanism they construct satisfies budget balance but will violate participation constraints at some histories. Bergemann and Valimaki [2010] build a dynamic version of the pivot mechanism that satisfies participation constraints but does not balance the budget budget.[7]

This paper is the most closely related to Athey and Miller [2007] and Skrzypacz and Toikka [2013]. Athey and Miller [2007] study the repeated bilateral trading problem under iid types, ex ante and ex post budget balance, and ex post incentive compatibility. They use a bounded budget account to show approximate efficiency under ex post budget balance. Skrzypacz and Toikka [2013] analyze the same problem with persistent types and multidimensional initial information. They establish a necessary and sufficient condition for efficiency under ex ante budget balance, thereby allowing for unbounded credit lines.[8]

While stressing voluntary participation in each period, we seek to characterize efficiency for an intermediate notion of budget balance, one that allows for the role of an intermediary with a bounded credit line. We also want to be able to impose greater restrictions on the cash flows to the intermediary. The implementation of the Collateral Dynamic VCG mechanism requires distribution of future surplus as collateral every period, in comparison to the one time participation fees in Skrzypacz and Toikka [2013], which may require large amounts of seed capital on the part of the agents, in addition to the unbounded credit line being offered by the mechanism designer.

[7]Under some additional assumptions, they also show that the mechanism satisfies an efficient exit condition– that is, the agents who stop being pivotal and also not assigned ay transfers.

[8]Using the balancing trick of Athey and Segal [2007, 2013], this condition also guarantees implementation under ex post budget balance, but then like Athey and Segal, a strong form of commitment is required on part of the agents by allowing individual rationality only in period 1.

In a companion paper, Lamba [2014], we generalize the main result presented here to an N-dynamic mechanism design problem. That paper also takes a closer look at the degree of transparency in the dynamic mechanisms, and their role in achieving desired objectives.

1.2 Model

Two agents, each with private information, agree to be in a dynamic bilateral trading relationship for a non-durable good. The buyer (B) has a hidden valuation for the good and the seller (S) is endowed with a technology to produce the good each period at a hidden cost. We assume that the buyer's valuation and the seller's cost are random variables[9], denoted v and c, distributed according to priors F and G on $\mathcal{V} = \{v_1, ..., v_N\}$ and $\mathcal{C} = \{c_1, ..., c_M\}$, that evolve according to independent Markov processes $F(.|.) : \mathcal{V} \times \mathcal{V} \to [0, 1]$ and $G(.|.) : \mathcal{C} \times \mathcal{C} \to [0, 1]$, respectively.[10] The densities have full support and are denoted by f, g, $f(.|.)$ and $g(.|.)$, respectively. Denote $\Delta v_{i+1} = v_{i+1} - v_i$, $\Delta c_{j+1} = c_{j+1} - c_j$. For the ease of notation, we will often write $\underline{v} = v_1$ and $\overline{c} = c_M$.

We choose the discrete type model for three reasons. First, it allows us to elucidate the key economic forces without measure theoretic complications. Second, it allows us to do comparative statics for simple examples, specifically when both agents have two possible types. And, finally many applications of dynamic mechanism design use numerical methods which require a discrete state space.

Each period p_t determines the probability of trade, that is, the production and allocation of the good from the seller to the buyer, $x_{B,t}$ the transfer from the buyer to the mechanism designer, and $x_{S,t}$ the transfer to the seller from the mechanism

[9]These shall be interchangeably referred to as their types.

[10]All the main results can accommodate moving supports. It would simply entail a change of notation to $\mathcal{V}_t$ and $\mathcal{C}_t$, to denote the respective supports in each period.

designer. The mechanism designer here can be considered as a financial intermediary, an institution as part of a larger social contract facilitating trade, or a simple transfer scheme in case $x_B = x_S$. The per period payoffs are given by $v_t p_t - x_{B,t}$ and $x_{S,t} - c_t p_t$ for the buyer and seller respectively.[11,12]

Taking the institutional details as given, both the buyer and seller can commit to the mechanism. The institutional details temper the role commitment will play in the model, as we elaborate below. Both the agents know their first period valuation and cost respectively when the contract is signed, and these then stochastically evolve over time. This assumption is crucial for it sows the seeds of asymmetric information in the model with commitment.

The (contractual) relationship lasts for T discrete periods, where $T \leq \infty$. Both the agents discount future payoffs with a common discount factor δ. The static version of this model, $\delta = 0$, with continuous type spaces is the one studied by Myerson and Satterthwaite [1983], and Chatterjee and Samuelson [1983]. For the the general discrete type space model, as pointed out by Myerson and Satterthwaite [1983], we can get possibility results for some measure of parameters. However as the number of types becomes large and the model converges to the continuous type space model, the measure of parameters for which efficiency can implemented converges to zero.

It is easy to show that a form of revelation principle holds and thus we can, without loss of generality, consider direct mechanisms. Every period the agents learn their own types, and then send a report to the mechanism, which in turn, spits out the allocation and transfers rules. Employing the revelation principle, however, demands a moral call on the information the mechanism itself reveals to the agents. In particular, does the buyer observe the seller's announcement and vice-versa? We

[11]The t subscript will not be used when the set of histories make the time dimension obvious.

[12]An equivalent model is one where the seller is endowed with a good every period and needs to decide whether she should sell the it to the buyer or consume it.

shall work in an environment where the announcements are publicly observed. There is a close information theoretic relationship between this *public* mechanism and the *blind* one where the announcements are not publicly observed. We refer the reader to Lamba [2014] for a discussion on this.[13]

The direct mechanism, say m, is then a collection of history dependent probability and transfer vectors, $m = \langle \mathbf{p}, \mathbf{x} \rangle = \left(p\left(v_t, c_t \left| h^{t-1}\right.\right), x_B\left(v_t, c_t \left| h^{t-1}\right.\right), x_S\left(v_t, c_t \left| h^{t-1}\right.\right)\right)_{t=1}^{T}$, where h^{t-1} and (v_t, c_t) are, respectively, the public history up to period $t-1$ and the types revealed at time t. These can also be succinctly written as $p(h^t)$, etc. In general, h^t is defined recursively as $h^t = \{h^{t-1}, (v_t, c_t)\}$, with $h^0 = \emptyset$. The set of possible histories at time t is denoted by H^t (for simplicity $H = H^T$).

The strategies of the buyer and the seller can potentially depend on a richer set of histories. For the buyer, the information available before his period t report is given by $h_B^t = \{h_B^{t-1}, v_{t-1}, \hat{v}_t\}$, where v_{t-1} is the announced type in period $t-1$, and $\hat{v}_t$ is the actual type in period t, starting with $h_B^0 = \{\hat{v}_1\}$. The seller's information is analogously defined. Let the set of private histories at time t be denoted by H_B^t and H_S^t, respectively. Thus, for a given mechanism, the strategy for the buyer, $(\sigma_{B,t})_{t=1}^{T}$, is then simply a function that maps private history into an announcement every period, $\sigma_{B,t} : H_B^t \mapsto \mathcal{V}$, and similarly for the seller, $\sigma_{S,t} : H_S^t \mapsto \mathcal{C}$.

1.3 The institutional framework

The edifice of the institutional machinery has three key foundations: private information, voluntary participation and limits on insurance. In the mechanism design lexicon, these would respectively be associated with incentive compatibility, individual rationality and budget balance constraints.

For a fixed mechanism m and strategies $\sigma = (\sigma_B, \sigma_S)$, the expected utilities on the

[13] In particular, permissibility of the results is an increasing function of the transparency of the mechanism.

induced allocation and transfers, after each possible history are defined as follows.

$$U_B^{m,\sigma}(h_B^t) = \mathbb{E}^{m,\sigma}\left[\sum_{\tau=t}^{T}\delta^{\tau-1}\left(v_\tau p_\tau - x_{B,\tau}\right)|h_B^t\right] \tag{1.1}$$

and,

$$U_S^{m,\sigma}(h_S^t) = \mathbb{E}^{m,\sigma}\left[\sum_{\tau=t}^{T}\delta^{\tau-1}\left(x_{S,\tau} - c_\tau p_\tau\right)|h_S^t\right] \tag{1.2}$$

Though, along truthful histories the difference between public and private histories is moot, and thus in much of what follows we shall suppress the same. Let $U_i^m = U_i^{m,\sigma^*}$, for $i = B, S$; where σ^* is the truth-telling strategy.

1.3.1 A change of variables

We propose a change of variables in the structure of the mechanism that will be central in our endeavor to establish a tight characterization of efficiency. In order to keep notation simple we suppress the type/variable over which expectation is taken. For example

$$p(v_t|h^{t-1}) = \sum_{j=1}^{M} p\left(v_t, c_{j,t}\left|h^{t-1}\right.\right) g(c_{j,t}|c_{t-1}),$$

where c_{t-1} is the $t-1$ period announcement of the seller, known to the buyer, and,

$$p(v_{t+1}|h^{t-1}, v_t) = \sum_{j=1}^{M}\sum_{k=1}^{M} p\left(v_{t+1}, c_{k,t+1}\left|h^{t-1}, v_t, c_{j,t}\right.\right) g(c_{j,t}|c_{t-1})g(c_{k,t+1}|c_{j,t})$$

Expected utility of the buyer can be recursively defined as[14]

$$U_B(v_t, c_t|h^{t-1}) = v_t p(v_t, c_t|h^{t-1}) - x_B(v_t, c_t|h^{t-1}) + \delta\sum_{i=1}^{N} U_B(v_{i,t+1}|h^{t-1}, v_t, c_t) f(v_{i,t+1}|v_t) \tag{1.3}$$

[14]For simplicity, when the mechanism m being employed is obvious, we simple write U_B and U_S suppressing the m.

and,

$$U_B(v_t|h^{t-1}) = v_t p(v_t|h^{t-1}) - x_B(v_t|h^{t-1}) + \delta \sum_{i=1}^{N} U_B(v_{i,t+1}|h^{t-1}, v_t) f(v_{i,t+1}|v_t)$$

Utility of the buyer of type v_t from misreporting (once) to be type v'_t, for a fixed type c_t of the seller, can be succinctly written as

$$U_B(v'_t; v_t, c_t|h^{t-1}) = v_t p(v'_t, c_t|h^{t-1}) - x_B(v'_t, c_t|h^{t-1}) +$$

$$\delta \sum_{i=1}^{N} U_B(v_{i,t+1}|h^{t-1}, v'_t, c_t) \cdot f(v_{i,t+1}|v_t)$$

$$= U_B(v'_t, c_t|h^{t-1}) + (v_t - v'_t) p(v'_t, c_t|h^{t-1}) +$$

$$\delta \sum_{i=1}^{N} U_B(v_{t+1,i}|h^{t-1}, v'_t, c_t) \cdot (f(v_{i,t+1}|v_t) - f(v_{i,t+1}|v'_t)) \tag{1.4}$$

Similarly,

$$U_B(v'_t; v_t|h^{t-1}) = U_B(v'_t|h^{t-1}) + (v_t - v'_t) p(v'_t|h^{t-1}) +$$

$$\delta \sum_{i=1}^{N} U_B(v_{i,t+1}|h^{t-1}, v'_t) \cdot (f(v_{i,t+1}|v_t) - f(v_{i,t+1}|v'_t))$$

The seller's utility, U_S, can be similarly defined.

It is straightforward to note that a mechanism $m = \langle \mathbf{p}, \mathbf{x} \rangle$, which is a collection of history dependent allocation and transfer vectors, can be equivalently defined to be $m = \langle \mathbf{p}, \mathbf{U} \rangle$, where (fixing the allocation) the duality between transfers and expected utility vectors is completely described by equation (2.1).

1.3.2 Incentive compatibility

Exploiting the one-deviation principle, incentive compatibility can be defined as follows.

Definition 1. *A mechanism $m = \langle \mathbf{p}, \mathbf{U} \rangle$ satisfies perfect Bayesian incentive compat-*

ibility if

$$U_B(v_t|h^{t-1}) \geq U_B(v_t'; v_t|h^{t-1}) \quad and \quad U_S(c_t|h^{t-1}) \geq U_S(c_t'; c_t|h^{t-1})$$

$\forall v_t, v_t' \in \mathcal{V}$, $\forall c_t, c_t' \in \mathcal{C}$, $\forall h^{t-1} \in H^{t-1}$, $\forall t$.

It states that along all truthful histories, buyer and the seller have no incentive to misreport their type.

1.3.3 Individual rationality

Even though commitment is assumed as part of our institutional architecture, we allow the agents to walk away after learning their type after any history if their utility from continuing in the contract falls below their reservation thresholds, which are normalized to zero. In keeping with our notation, we have:

Definition 2. *A mechanism $m = \langle \mathbf{p}, \mathbf{U} \rangle$ satisfies perfect Bayesian individually rationality if*

$$U_B(v_t|h^{t-1}) \geq 0 \quad and \quad U_S(c_t|h^{t-1}) \geq 0$$

$\forall v_t \in \mathcal{V}$, $\forall c_t \in \mathcal{C}$, $\forall h^{t-1} \in H^{t-1}$, $\forall t$.

We say that a mechanism is perfect Bayesian implementable if it is perfect Bayesian incentive compatible and individually rational.

1.3.4 Budget balance

In mechanism design with many agents budget balance is seen as the limits on insurance or external subsidies available to them. In addition to the traditional notions of ex ante and ex post budget balance, we introduce an intermediate notion of interim budget balance.

We say that a mechanism is *interim budget balanced* if

$$\mathbb{E}^m\left[\sum_{\tau=t}^{T}\delta^{\tau-t}\left(x_{B,\tau}-x_{S,\tau}\right)\mid h^{t-1}\right]\geq 0$$

$\forall\, h^{t-1}\in H^{t-1}$.[15] The mechanism is *ex ante budget balanced* if interim budget balance holds for the null history. Moreover, we say that the mechanism is *ex post budget balanced* if the entire vector of transfers are equal for any history, $x_B = x_S$.

Next, using equations (1.1) and (1.2) we can write the expected budget surplus that a mechanism generates after any history h^{t-1} to be

$$EBS(h^{t-1})=\mathbb{E}^m\left[\sum_{\tau=t}^{T}\delta^{\tau-t}\left(v_\tau-c_\tau\right)p_\tau-U_B(v_t|h^{t-1})-U_S(c_t|h^{t-1})\mid h^{t-1}\right] \quad (1.5)$$

The ex ante budget surplus is denoted simply by $EBS = EBS(h^0)$. We have,

Definition 3. *A mechanism $\langle \mathbf{p}, \mathbf{U}\rangle$ satisfies ex ante budget balance if*

$$EBS\geq 0$$

This is the weakest possible notion of budget balance for this dynamic model. It means that the mechanism designer does not loose money in an expected ex ante sense. A more robust definition of budget balance in our opinion, which still allows for the role of an intermediary is the one where a positive budget surplus is guaranteed after every history.

Definition 4. *A mechanism $\langle \mathbf{p}, \mathbf{U}\rangle$ satisfies interim budget balance if*

$$EBS(h^{t-1})\geq 0 \quad \forall h^{t-1}\in H^{t-1},\ \forall t$$

[15] The exact definition will employ almost sure notions on the set of histories. It will be obvious and is suppressed for the ease of exposition.

This can be motivated in many ways. First, it can be viewed as a participation constraint for the mechanism designer- after any history, just like the the two agents, the mechanism designer must have an incentive to continue in the relationship. Second, it is a bankruptcy constraint for the intermediary. If the contract reaches a stage the where the intermediary is expected to loose money, he or she should be allowed to shut shop.

Interim budget balance does not allow the mechanism designer to draw from past surplus. It can be viewed as a constraint on the intermediary's commitment power. While this would a reasonable assumption in many contexts and an interesting benchmark in it own right, it is also important to note that the family of constraints defining interim budget balance can easily be generalized to a class where the mechanism designer is allowed to save. We take up this issue in section 1.8.

Finally, the most standard (and strictest) definition of budget balance from the static literature that can be generalized to dynamic environments states the transfers should exactly equal across all histories for all time periods.

Definition 5. *A mechanism $m = \langle \mathbf{p}, \mathbf{x} \rangle$ satisfies ex post budget balance if*

$$x_B(v_t, c_t|h^{t-1}) - x_S(v_t, c_t|h^{t-1}) = 0,$$

$\forall v_t \in \mathcal{V}$, $\forall c_t \in \mathcal{C}$, $\forall h^{t-1} \in H^{t-1}$, $\forall t$.[16]

A natural way to motivate this in the dynamic model is the absence of an outside

[16]Equivalently, a mechanism $m = \langle \mathbf{p}, \mathbf{U} \rangle$ satisfies ex post budget balance if

$$(v_t - c_t)p(v_t, c_t|h^{t-1}) - \left(U_B(v_t, c_t|h^{t-1}) - \delta \sum_{i=1}^{N} U_B(v_{t+1,i}|h^{t-1}, v_t, c_t) f(v_{t+1,i}|v_t) \right) -$$

$$\left(U_S(v_t, c_t|h^{t-1}) - \delta \sum_{j=1}^{M} U_S(c_{t+1,j}|h^{t-1}, v_t, c_t) g(c_{t+1,j}|c_t) \right) = 0$$

$\forall v_t \in \mathcal{V}$, $\forall c_t \in \mathcal{C}$, $\forall h^{t-1} \in H^{t-1}$, $\forall t$.

insurance provider or financial intermediary. A contractual relationship is thus more the order of interpretation rather than a mechanism. Both the agents insure each other against bad shocks, and premium paid can be recovered through continuation utility.

Note the hierarchy in budget balance

$$\text{ex post budget balance } \Rightarrow \text{ interim budget balance} \Rightarrow \text{ ex ante budget balance}$$

It is reasonably easy to show, by (backward) inductive redistribution of transfers and replication of an argument from static mechanism design for every period[17], that if T is finite and there exists a mechanism that implements the efficient allocation under interim budget balance, then there also exists a mechanism that implements it under ex post budget balance. We show in section 1.5 that a similar result can be constructed when $T = \infty$. However, it will require large (but bounded) payments on the part of the agents after certain histories.

1.3.5 Objectives

One of the most widely accepted objectives of mechanism design is that of efficiency.[18] We shall invoke the strongest possible version in its ex post form.

Definition 6. *A mechanism $m = \langle \mathbf{p}, \mathbf{U} \rangle$ satisfies efficiency if*

$$p(v_t, c_t | h^{t-1}) = \begin{cases} 1 & \textit{if } v_t > c_t \\ 0 & \textit{otherwise} \end{cases}$$

$\forall v_t \in \mathcal{V}$, $\forall c_t \in \mathcal{C}$, $\forall h^{t-1} \in H^{t-1}$, $\forall t$.

Thus, regardless of history, under a positive instantaneous surplus and only then,

[17]See Lemma 3 in the appendix.

[18]See Holmstrom and Myerson [1983] for the various notions of efficiency.

efficiency demands trade and always with probability 1. This is a straightforward generalization of the notion typically used in static models, and allows for a direct comparison with Myerson and Satterthwaite [1983].

1.3.6 The interaction of constraints

Before we jump into the results, it worthwhile to investigate the interaction of the various forces laid out in this section. From static mechanism design we know that it is the interaction of private information (and hence incentive compatibility) with participation (hence individual rationality) and budget balance that leads to impossibility of efficiency result of Myerson and Satterthwaite [1983]. If we wished to implement the efficient allocation under only individual rationality, the well known Vickery-Clarkes-Groves (VCG) mechanism does so. On the other hand, if we wished to implement the efficient allocation under only budget balance, the d'Aspremont and Gérard-Varet (AGV) mechanism does so. It is the simultaneous interaction of the three forces that leads to a departure from efficiency.

So, in order to understand how dynamics can overcome the impossibility results, we must model the simultaneous interactions of these forces every period. While ex ante budget balance is a good benchmark to have, it only requires budget balance and individual rationality to interact initially, thereby the possibility results are not *exclusively* because of the dynamics of the problem. Also, it may not be a plausible restriction for many real contractual situations because it forces the intermediary to subsidize trade by unbounded amounts along some histories as time horizon gets long. On the other hand, ex post budget balance might be too restrictive when agents can contract dynamically. Our endeavor in this paper is to model the simultaneous interaction of the three forces by requiring a continuation budget surplus in expectation every period.

1.4 A simple model

To fix ideas we first consider the simple exercise of implementing the efficient allocation in a two period model where both the agents have two possible types. So, in terms of the model described in section 2: $T = 2$, $N = M = 2$. In addition, we assume that the agents' types are arranged in the following order: $v_H > c_H > v_L > c_L$. Any other ordering will render posted prices efficient statically and hence dynamically.

Consider the following intuitive pricing mechanism. Each period the buyer announces (suggests) a price P_B that he may pay for the object, and the seller similarly announces a price P_S she may accept for producing and selling it. Based on these announcements, trade takes place if $P_B \geq P_S$, that is $\hat{p}(P_B, P_S) = \mathbf{1}_{\{P_B \geq P_S\}}$; and the transfers are defined as follows:

$$\hat{x}_B(P_B, P_S) = \begin{cases} \min\{v_k | v_k \geq P_S\} & \text{if } P_B \geq P_S \\ 0 & \text{otherwise} \end{cases}$$

$$\hat{x}_S(P_B, P_S) = \begin{cases} \max\{c_k | P_B \geq c_k\} & \text{if } P_B \geq P_S \\ 0 & \text{otherwise} \end{cases}$$

Call the above mechanism $\hat{m}$, and note that it is history independent. It is straightforward to show that announcing the true valuation every period, $P_B(v_i) = v_i, P_S(c_j) = c_j$, is a weakly dominant strategy for both the agents. It follows from the fact that $\hat{m}$ is essentially a discrete counterpart to the famous VCG mechanism[19]. The added structure is provided to ensure it is statically the minimalist efficient mechanism in terms of the information rent paid to the agents[20], and maximalist in terms of the surplus provided to the mechanism designer. Without loss of generality,

[19]Since this is a two player mechanism design problem, in the VCG mechanism each player would be asked to pay the externality he/she imposes on the other player: $x_B^{vcg}(P_B, P_S) = P_S \mathbf{1}_{\{P_B \geq P_S\}}$ and $x_S^{vcg}(P_B, P_S) = P_B \mathbf{1}_{\{P_B \geq P_S\}}$.

[20]As will be shown in the general environment later, this mechanism ensures that the local incentives hold as equalities.

in what follows, we restrict the announcement space of the buyer and seller to be $\{v_H, v_L\}$ and $\{c_H, c_L\}$, respectively. Then,

$$\hat{x}_B(v_H, c_H) = v_H, \hat{x}_B(v_H, c_L) = v_L, \hat{x}_B(v_L, c_H) = 0, \hat{x}_B(v_L, c_L) = v_L$$

$$\hat{x}_S(v_H, c_H) = c_H, \hat{x}_S(v_H, c_L) = c_H, \hat{x}_S(v_L, c_H) = 0, \hat{x}_S(v_L, c_L) = c_L$$

The obvious problem with this mechanism is that it may a run a deficit, and budget balance will be violated even in an expected sense. Thus, more cash needs to flow to the mechanism designer. So as a final step, in addition to the above efficient mechanism, we allow both players to deposit some additional money with the mechanism designer every period. Dynamics present this possibility in the form of expected future economic surplus. Extraction of this additional remuneration for the mechanism designer must respect the agents' incentive compatibility and individual rationality constraints.

Suppose r_1^B and $r_2^B(h)$ are the costs paid by the buyer in period 1 and in period 2 after history h, respectively. The local incentive constraints hold as equalities in m, so more information rent cannot be drawn back from any type. Thus, an intuitive way to construct these payments is extracting from each type of the buyer the expected rent earned by the "lowest" type in excess of his reservation value of zero. Define

$$r_2^B(h) = U_B^{\hat{m}}(v_L|h) \;\forall\; h \in \{v_H, v_L\} \times \{c_H, c_L\}, \text{ and}$$

$$r_1^B + \delta \mathbb{E}^{\hat{m}} \left[r_2^B(\tilde{h}) \right] = U_B^m(v_L)$$

A back of the envelope calculation shows that $U_B^{\hat{m}}(v_L|h) = 0$ for all h, and

$$U_B^{\hat{m}}(v_L) = \delta f(v_H|v_L) \left[g(c_H) g(c_L|c_H) + g(c_L) g(c_L|c_L) \right] \Delta v$$

The costs paid by the seller in the two periods: r_1^S and $r_2^S(h)$, can be similarly defined.

Call this new mechanism m^*: $p_t^*(v_i, c_j) = \mathbf{1}_{\{v_i \geq c_j\}}$ for $t = 1, 2$ independent of history, and

$$x_{B,1}^*(v_i, c_j) = \hat{x}_B(v_i, c_j) + r_1^B, \quad x_{B,2}^*(v_i, c_j|h) = \hat{x}_B(v_i, c_j) + r_2^B(h),$$

$$x_{S,1}^*(v_i, c_j) = \hat{x}_S(v_i, c_j) - r_1^S, \quad x_{S,2}^*(v_i, c_j|h) = \hat{x}_S(v_i, c_j) - r_2^S(h)$$

In mechanism m^*, the incentive constraints hold tightly and the participation constraints of the “lowest" type bind. Thus, given informational and participation constraints, no more rents can be extracted from the agents. If the cash flow to the mechanism designer in m^* does not create an expected budget surplus, no other mechanism will.

Define the ex ante budget surplus for this two period model to be

$$EBS = \sum_{i=H,L} \sum_{j=H,L} f(v_i)g(c_j) \Big[x_B(v_i, c_j) - x_S(v_i, c_j) +$$

$$\delta \sum_{k=H,L} \sum_{l=H,L} f(v_k|v_i)g(c_l|c_j) \left(x_B(v_k, c_l|v_i, c_j) - x_S(v_k, c_l|v_i, c_j) \right) \Big]$$

and the expected surplus in period 2 after history $h = (h_v, h_c)$ to be

$$EBS(h) = \sum_{k=H,L} \sum_{l=H,L} f(v_k|h_v)g(c_l|h_c) \Big[x_B(v_k, c_l|h) - x_S(v_k, c_l|h) \Big]$$

Letting EBS^* and $EBS^*(h)$ denote the above entities for the mechanism m^*, we can completely characterize the implementation of the efficient allocation for our simple model by the following result.

Proposition 1. *There exists an incentive compatible and individually rational mechanism that implements the efficient allocation under interim budget balance if and only if* $EBS^* \geq 0$, *and* $EBS^*(h) \geq 0$ $\forall h$.

This result pins down the necessary and sufficient conditions for the implementation of the efficient allocation under interim budget balance in our simple two period model. It allows budget to balanced dynamically, but with an added “participation" constraint for the mechanism designer at every history. If budget was only required to be balanced ex ante, the above mechanism and result still go through with the necessary and sufficient condition being simply $EBS^* \geq 0$.

A plausible alternative to our mechanism could be one in which the agents deposit money with the mechanism designer only at the beginning of the contract, satisfying only ex ante budget balance, and draw upon it whenever the need arise.[21] While this may certainly be a good assumption in various contractual scenarios, in the general model with many type and time periods, there are at least two situations in which this may not be an appropriate approximation of a real dynamic contract. First, the agents may not have large amounts of seed capital to deposit from the word go. Second, the mechanism designer designer would be asked to subsidize trade by arbitrarily large amounts under some histories. at which stage he can possibly file for bankruptcy.

In contrast the mechanism presented here requires the agents to deposit a smaller fees every period post which (whenever possible) a simple pricing mechanism implements the efficient allocation.[22] Fixing the efficient allocation as the objective, and given incentive and participation constraints, it precisely pins down the maximal possible continuation surplus that can be extracted from the economic relationship of the buyer and the seller in expectation at every history.

What about implementation under ex post budget balance? If $\delta = 0$, that is we were in a static world, it has been documented by a series of papers that if there

[21]Skrzypacz and Toikka [2013] construct such a mechanism.

[22]In the simple model, since the deposit in period 2 is zero, that is $r_2^B(h) = r_2^S(h) = 0 \ \forall h$, the two mechanisms coincide. But it'll be clear in sections 1.5 and 1.6 that in the general model deposits after period 1 are typically non-zero.

exists an incentive compatible and individually rational mechanism (p, x_B, x_S) that satisfies $EBS \geq 0$, then there also exists a mechanism $(p, \tilde{x}_B, \tilde{x}_S)$ implements p under ex post budget balance; so $\tilde{x}_B = \tilde{x}_S$.[23] The intuition for this result is simply that if ex ante budget surplus is generated by a mechanism, and we only need to satisfy Bayesian incentive compatibility (as opposed to dominant strategy), then this surplus can be re-distributed between the agents across types to construct a mechanism that implements the same allocation under ex post budget balance.

For the dynamic model, define

$$EBS_1 = \sum_{i=H,L} \sum_{j=H,L} f(v_i) g(c_j) \left[x_B(v_i, c_j) - x_S(v_i, c_j) \right]$$

and $EBS_2(h) = EBS(h)$ to be the *current* expected budget surplus generated each period.[24] In a direct generalization of the static result described above we can show that there exists a mechanism that implements an allocation p under ex-post budget balance if and only if it can be implemented by a mechanism that satisfies $EBS_1 \geq 0$ and $EBS_2(h) \geq 0$.[25]

Suppose $EBS^* \geq 0$ and $EBS^*(h) \geq 0$ for all h. Then, the only way efficient allocation cannot be implemented under ex post budget balance is if $EBS_1^* < 0$. In that case we can simply move transfers across periods in an incentive compatible and individually rational manner to generate a new mechanism, say m^{**}, that satisfies $EBS^{**}(h) = 0$ for all h, and $EBS_1^{**} \geq 0$.

Fix some $\alpha \in [0, 1]$. Define

$$x_{B,2}^{**}(v_i, c_j | h) = x_{B,2}^{*}(v_i, c_j | h) - \alpha EBS^*(h)$$

[23] See Mailath and Postlewaite [1991].

[24] Note that $EBS = EBS_1 + \delta \mathbb{E} \left[EBS_2(\tilde{h}) \right]$.

[25] Lemma 3 in the appendix.

$$x^{**}_{B,1}(v_i, c_j) = x^{*}_{B,1}(v_i, c_j) + \delta\alpha EBS^{*}(v_i, c_j)$$

$$x^{**}_{S,2}(v_i, c_j|h) = x^{*}_{S,2}(v_i, c_j|h) + (1-\alpha)EBS^{*}(h)$$

$$x^{**}_{S,1}(v_i, c_j) = x^{*}_{S,1}(v_i, c_j) - \delta(1-\alpha)EBS^{*}(v_i, c_j)$$

Thus, we have the characterization of implementation of the efficient allocation under ex post budget balance.

Corollary 1. *There exists an incentive compatible and individually rational mechanism that implements the efficient allocation under ex post budget balance if and only if* $EBS^{*} \geq 0$, *and* $EBS^{*}(h) \geq 0 \; \forall h$.

It is clear that in the absence of any additional constraints the necessary and sufficient condition for implementation of the efficient allocation under interim and ex post budget balance are exactly the same. Thus, in view of Corollary 1, should interim and ex post budget balance be considered “equivalent"?

Ex post budget balance in the dynamic framework allows for the possibility of the agents insuring each other against bad shocks through direct transfers. Interim budget balance on the other hand is laxer, and allows for the role of an intermediary. Since types are imperfectly correlated, a long enough time horizon (and high enough discounting) will allow the possibility of such promises to be kept. However, for high levels of persistence of bad shocks, the transfers required to meet these promises will be arbitrarily large.

Note that $x^{**}_{B,1}(v_i, c_j) \geq x^{*}_{B,1}(v_i, c_j)$ and $x^{**}_{S,1}(v_i, c_j) \leq x^{*}_{S,1}(v_i, c_j)$ with at least one strict inequality. Thus, in the presence of a hard (upper) bound on per-period transfers it is easier to satisfy interim budget balance than ex post. These bounds, which can be reasonably considered to be credit constraints, create the role for an intermediary who can break the budget inter-temporally. We will have to say on this in section 1.8.

A second-best formulation. An alternative way to look at the same problem is through the prism of maximizing gains from trade or surplus. Using the revelation principle, the problem can simply be stated as

$$\max_{\langle \mathbf{p},\mathbf{x}\rangle} \quad \mathcal{S} = \sum_{i=H,L}\sum_{j=H,L} f(v_i)g(c_j)\Big[(v_i - c_j)\, p(v_i, c_j) +$$

$$\delta \sum_{k=H,L}\sum_{l=H,L} f(v_k|v_i)g(c_l|c_j)\,(v_k - c_l)\, p(v_k, c_l|v_i, c_j)\Big]$$

subject to

$$IC_H^a, IC_L^a, IR_H^a, IR_L^a, \qquad \text{for } a = B, S$$

$$IC_H^a(h), IC_L^a(h), IR_H^a(h), IR_L^a(h), \qquad \text{for } a = B, S \text{ and } h \in \{v_H, v_L\} \times \{c_H, c_L\}$$

and,

$$EBS = \sum_{i=H,L}\sum_{j=H,L} f(v_i)g(c_j)\Big[x_B(v_i, c_j) - x_S(v_i, c_j) +$$

$$\delta \sum_{k=H,L}\sum_{l=H,L} f(v_k|v_i)g(c_l|c_j)\,(x_B(v_k, c_l|v_i, c_j) - x_S(v_k, c_l|v_i, c_j))\Big] \geq 0$$

$$EBS(h) = \sum_{k=H,L}\sum_{l=H,L} f(v_k|h_v)g(c_l|h_c)\Big[x_B(v_k, c_l|h) - x_S(v_k, c_l|h)\Big] \geq 0$$

$$\text{for } h = (h_v, h_c) \in \{v_H, v_L\} \times \{c_H, c_L\}$$

where IC_i^a and IR_i^a are respectively the incentive compatibility and individual rationality constraint in period 1 for agent a of type i, $IC_i^a(h)$ and $IR_i^a(h)$ are analogously the constraints in period 2 after history h, and $EBS \geq 0$ and $EBS(h) \geq 0$ represent the interim budget balance constraints in period 1 and period 2 after history h respectively. Also, note that if we take away the four constraints: $EBS(h) \geq 0$, and require the budget to balance only ex ante, we would of course get more permissible results on efficiency.

The choice of surplus as the maximand, as opposed to any other point on the Pareto frontier, is driven by the fact that whenever the parameters of the problem allow the efficient allocation to be implemented, it will indeed be the unique solution of this second-best.

A natural way to proceed with this problem is to consider a relaxed problem with "local downward" incentive constraints: $IC_H^B, IC_L^S, IC_H^B(h), IC_L^S(h)$, and the participation constraint of the "lowest" type: $IR_L^B, IR_H^S, IR_L^B(h), IR_H^S(h)$, and the (five) interim budget balance constraints. It is fairly straightforward to show that the solution to this relaxed problem is indeed the solution to the original one.[26]

Next, the highest values of EBS and $EBS(h)$ are obtained when the incentive compatibility and individual rationality constraints in the relaxed problem hold as equalities. Using these equalities, the transfers can be eliminated from EBS and $EBS(h)$, which can be re-written in terms of the allocation p. We consider the problem of the second-best first under ex ante and then interim budget balance.

Now, define the ex ante virtual valuations, that is the valuations net of the information rents, associated with each possible realization and history of types.

$$MR(v_H) = v_H, MR(v_L) = v_L - \frac{f(v_H)}{f(v_L)}\Delta v, MC(c_L) = c_L, MC(c_H) = c_H + \frac{g(c_L)}{g(c_H)}\Delta c$$

$$MR(v_i|v_H) = v_i, \text{ for } i = H, L, \text{ and } MC(c_j|c_L) = c_j \text{ for } j = H, L$$

$$MR(v_H|v_L) = v_H, MR(v_L|v_L) = v_L - \frac{f(v_H)}{f(v_L)}\frac{f(v_L|v_L) - f(v_L|v_H)}{f(v_L|v_L)}\Delta v$$

$$MC(c_L|c_H) = c_L, MC(c_H|c_H) = c_H - \frac{g(c_L)}{g(c_H)}\frac{g(c_H|c_H) - g(c_H|c_L)}{g(c_H|c_H)}\Delta c$$

The notation MR and MC is in the spirit of Bulow and Roberts [1983], where the virtual valuations of the buyer and the seller are motivated as the marginal

[26] Note that this is possible because of two types assumption. For the second-best, in a dynamic contract with more than two types global constraints typically bind. See Battaglini and Lamba [2014].

revenue and marginal cost respectively, lending a Ramsey pricing interpretation to the mechanism design problem of bilateral trading. Also, note that for both agents, distortions persist in the second period only when their type is “low" in period 1.[27]

Define the ex ante virtual surplus to be

$$EBS(p) = \sum_{i=H,L}\sum_{j=H,L} f(v_i)g(c_j)\Big[\left(MR(v_i) - MC(c_j)\right)p(v_i, c_j) +$$

$$\delta \sum_{k=H,L}\sum_{l=H,L} f(v_k|v_i)g(c_l|c_j)\left(MR(v_k|v_i) - MC(c_l|c_j)\right)p(v_k, c_l|v_i, c_j)\Big]$$

The ex ante budget constraint is then simply $EBS(p) \geq 0$. If the budget was only required to balance in an ex ante sense, transfers can be moved across time and types freely as long they respect the the individual rationality of the agents and the ex ante budget surplus condition. Thus, maximization of surplus $\mathcal{S}$ with respect to the virtual surplus constraint being positive, $EBS(p) \geq 0$ is just an extension of the static Ramsey pricing problem.

All types (in period 1 and 2) for which $MR > MC$, trade will occur. For the rest, a ranking based on the efficiency-profit ratio and the binding virtual surplus constraint determine (im)possibility of trade.[28] Note that ranking is homogenous for types in both periods, that is, for $\Pi(v_i, c_j) = MR(v_i) - MC(c_j) < 0$ and $\Pi(v_k, c_l|v_i, c_j) = MR(v_k|v_i) - MC(c_l|c_j) < 0$, trade is allowed in decreasing order of $(v - c)/(-\Pi)$ across periods till $EBS(p) \geq 0$ binds. It is important to note that more trade happens in period 1 of the dynamic problem than the static one, because a future (virtual) surplus relaxes the $EBS(p) \geq 0$ constraint. Finally, the efficient allocation

[27]This a product of the “generalized no distortion at the top" principle in dynamic contracts. Once the type process hits the “highest" type, distortions disappear and the virtual valuation is equal to the actual valuation for that agent.

[28]Since this a discrete type model, randomization occurs at the optimum for some types with $MR < MC$. It is easy to show (Myerson [1985]) that in the continuous type model the second-best is always bang bang: probability of trade is either zero or one.

can be implemented when ex ante budget surplus evaluated at the efficient allocation, denoted say by $EBS(p^*)$, is non-negative.

But, what if the expected budget has to be balanced in every period? Then, in addition to the ex ante budget balance constraint, the mechanism must also satisfy $EBS(h) \geq 0$ for all h. The *interim* virtual valuations can be defined as

$$MR_2(v_H|v_i) = v_H, MR_2(v_L|v_i) = v_L - \frac{f(v_H|v_i)}{f(v_L|v_i)}\Delta v$$

$$MC_2(c_L|c_j) = c_L, MC_2(c_H|c_j) = c_H + \frac{g(c_L|c_j)}{g(c_H|c_j)}\Delta c$$

Define the interim virtual valuation after history $h = (v_i, c_j)$ to be

$$EBS(p)(v_i, c_j) = \sum_{k=H,L}\sum_{l=H,L} f(v_k|v_i)g(c_l|c_j)\left[MR_2(v_k|v_i) - MC_2(c_l|c_j)\right]p(v_k, c_l|v_i, c_j)$$

It is easy to see that the interim budget balance constraint after history h is $EBS(p)(h)$ ≥ 0. The original problem of the second-best can then be restated simply as the maximization of $\mathcal{S}$ under $EBS(p) \geq 0$ and $EBS(p)(h) \geq 0$ for all h.

In contrast to the static model and the dynamic model under ex ante budget balance, the Ramsey pricing problem for period 2 has to internalize two types of marginal revenue and cost functions. In the first period, all types for which $MR > MC$ trade for sure. In the second period, all types for which $MR > MC$ and $MR_2 > MC_2$ trade for sure. For all other realizations, the *correct* virtual valuation for the Ramsey ranking contest in period 2 depends on whether constraint $EBS(p)(h) \geq 0$ binds. If after history $h = (v_i, c_j)$, it binds, sorting is done based on $(v - c)/(-\Pi_2)$, where $\Pi_2(v_k, c_l|v_i, c_j) = MR_2(v_k|v_i) - MC_2(c_l|c_j)$. Once we have run through all histories for which $EBS(p)(h) \geq 0$ binds, these allocations are substituted back in to equation $EBS(p) \geq 0$. Then, a Ramsey ranking is done for the remaining second period types and first period types as before.

In this relaxed problem if none of the budget balance conditions bind at the optimum, optimization yields the efficient allocation. Letting $EBS(p^*)$ and $EBS(p^*)(h)$ denote the virtual valuations evaluated at the efficient allocation, we get following result

Proposition 2. *There exists an incentive compatible and individually rational mechanism that implements the efficient allocation under interim budget balance if and only if* $EBS(p^*) \geq 0$, *and* $EBS(p^*)(h) \geq 0 \ \forall h$.

As described above, implementation under interim budget balance can be used a foundation to construct a new mechanism that satisfies ex post budget balance.

Corollary 2. *There exists an incentive compatible and individually rational mechanism that implements the efficient allocation under interim budget balance if and only if* $EBS(p^*) \geq 0$, *and* $EBS(p^*)(h) \geq 0 \ \forall h$.

It is of course easy to see that Propositions 1 and 2, provide the same necessary and sufficient conditions on the implementability of the efficient allocation, using separate constructive arguments.

Comparative Statics. What do Propositions 1 and 2 mean in terms of the parameters of the problem? It is clear that they put joint restrictions on the distance between types, level of discounting and the Markov matrix. In the rest of this section, we try to understand what moving parts mean in terms of economics of the problem.

First it easy to note that at $\delta = 0$ or with perfectly persistent types, efficiency holds if and only if

$$EBS(p^*)(\delta = 0) = \sum_{i=H,L} \sum_{j=H,L} f(v_i)g(c_j)\left[MR(v_i) - MC(c_j)\right] p(v_i, c_j) \geq 0$$

This is equivalent to the condition in Matsuo [1989] for the two types static bilateral trading problem. Note that this condition does hold for a significant measure of

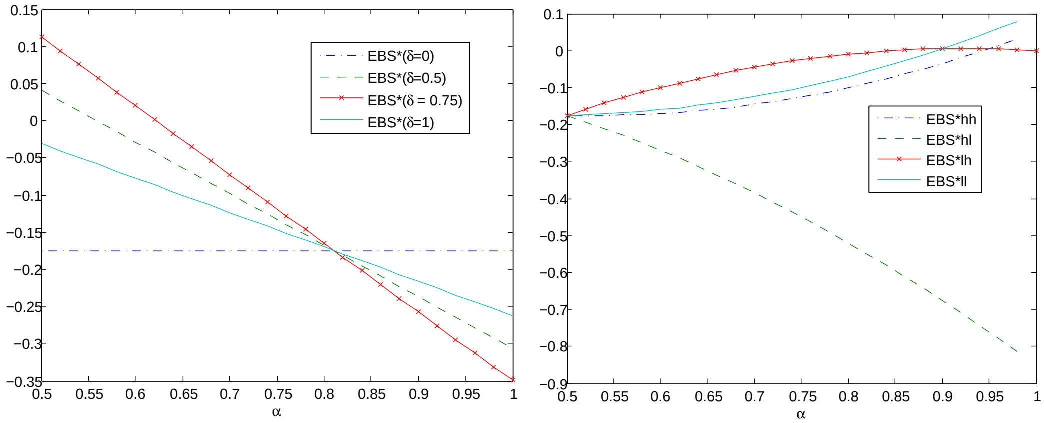

(a) EBS^{*} for $\delta = 0,\ 0.5,\ 0.75,\ 1$ (b) $EBS^{*}(v_i, c_j)$ for $i, j = H, L$

Figure 1.1: Expected Budget Surplus for two types, two periods model

parameters, unlike the continuous type space model where we get an impossibility of efficiency as long as there is any overlap in buyer and seller supports.[29]

We parametrize the problem for simple comparative statics that can be communicated through pictures. Let the prior for agents be uniform, and the transition probabilities be given by $f(v|v) = g(c|c) = \alpha$, $f(v|v') = g(c|c') = 1 - \alpha$ where $\alpha > \frac{1}{2}$. So α measures the level of persistence in the model. The type spaces- $\{v_H, v_L\}$ and $\{c_H, c_L\}$- are chosen such that $EBS(p^{*})(\delta = 0) < 0$; therefore, the possibility of efficiency is *exclusively* due to the dynamics of the problem.

Figure 1.1a plots EBS^{*} against α for four different values of δ. The horizontal line is the of course $EBS^{*}(\delta = 0)$. As is intuitive the ex ante budget surplus is a decreasing function of the level of persistence in the model, and an increasing function of discounting in the “positive" region of the economic surplus.

Figure 1.1b plots $EBS^{*}(v_i, c_j)$ for $i, j = H, L$. First note that all curves start at the same point,. This is because at $\alpha = 1/2$ the model is iid and history does not matter for the interim budget balance constraint. Next, this graph shows why it is not straightforward to rank expected budget surplus according to the level of

[29]However, for the general discrete type space model, as the number of types becomes large and the model converges to the continuous type space model, the measure of parameters where the condition holds converges to zero.

persistence in the model. The parameter α affects both the probability of types and level of distortions. And, as shown in the construction of interim virtual valuations, these distortions can go both ways as a function of α. The imperative thing to note though is that the minimum all of $EBS^*(v_i, c_j)$ is a decreasing function of α. So, persistence continues to be bad news even for the interim budget balance constraints.

1.5 Dynamic collateral VCG mechanism

Now, we ask the main question: *when can efficiency be sustained in a dynamic bargaining problem with Markovian private information under voluntary participation and (interim or ex post) budget balance?*

In order to answer this question we construct an incentive compatible mechanism that produces the minimal possible rent for the agents after any history and allocates maximal possible surplus to the mechanism designer. Borrowing techniques from Williams [1999] and Krishna and Perry [2000], we start with a VCG mechanism and adapt it to the discrete types framework to create a tight mechanism that satisfies local incentive constraints with equality.[30]

The standard VCG mechanism consists of $x_B^{vcg}(v_t, c_t|h^{t-1}) = c_t$, and $x_S^{vcg}(v_t, c_t|h^{t-1}) = v_t$, that is each agent pays the externality he/she imposes on the other agent. Since payoff equivalence does not "automatically" hold for discrete types, the gap between the types needs to taken care of in order to create a tight mechanism does not give rent to agents that is more than necessary to satisfy incentives. The following modified

[30] As presented in the simple model, an alternative approach could be to write down the problem of the second-best and back out conditions under which it implements the efficient allocation. We use the dynamic generalization of the VCG to allow for easier comparisons to the static literature and also to provide a simple way of implementing the efficient allocation.

version of the VCG mechanism takes care of the same.[31]

$$x_B^{Mvcg}(v_i, c_j|h^{t-1}) = \begin{cases} \min\{v_k|v_k > c_j\} & \text{if } v_i > c_j \\ 0 & \text{otherwise} \end{cases}$$

$$x_S^{Mvcg}(v_i, c_j|h^{t-1}) = \begin{cases} \max\{c_k|v_i > c_k\} & \text{if } v_i > c_j \\ 0 & \text{otherwise} \end{cases}$$

where $Mvcg$ stands for modified version of VCG.

The key contribution of this mechanism over the standard VCG is that it ensures the "downward" local constraints hold as equalities.

Lemma 1. *For all* $i < N$

$$U_B^{Mvcg}(v_{i+1}|h^{t-1}) = U_B^{Mvcg}(v_i|h^{t-1}) + \Delta v_{i+1} p(v_i|h^{t-1}) + \delta \sum_{k=1}^{N} U_B^{Mvcg}(v_k|h^{t-1}, v_i) \cdot (f(v_k|v_{i+1}) - f(v_k|v_i))$$

and for all $j > 0$

$$U_S^{Mvcg}(c_j|h^{t-1}) = U_S^{Mvcg}(c_{j+1}|h^{t-1}) + \Delta c_{j+1} p(c_{j+1}|h^{t-1}) + \delta \sum_{k=1}^{N} U_S^{Mvcg}(c_k|h^{t-1}, c_{j+1}) \cdot (g(c_k|c_j) - g(c_k|c_{j+1}))$$

This mechanism may or may not run a deficit. However, as the number of possible types becomes large (and the distance between them goes to zero), the mechanism converges to the standard VCG mechanism which we know violates any notion of budget balance. The challenge then is, using the modified VCG mechanism as a base, can we construct a mechanism that while preserving incentives and participation

[31]A static version of it appears in Manea and Kos [2009].

satisfies interim budget balance whenever the parameters of the problem deem it possible?

The following partial payoff equivalence result paves our way.

Lemma 2. *Suppose* $\langle \mathbf{p}, \mathbf{U} \rangle$ *is incentive compatible. For a family of finite constants* $(a_B(c_t|h^{t-1}), a_S(v_t|h^{t-1}))$*, define*

$$\tilde{U}_B(v_t, c_t|h^{t-1}) = U_B(v_t, c_t|h^{t-1}) + a_B(c_t|h^{t-1}), \text{ and}$$

$$\tilde{U}_S(v_t, c_t|h^{t-1}) = U_S(v_t, c_t|h^{t-1}) + a_S(v_t|h^{t-1})$$

Then $\left\langle \mathbf{p}, \tilde{\mathbf{U}} \right\rangle$ *is also incentive compatible.*

This result provides a potent history dependent formulation of the payoff equivalence result. It provides a structure to the composition of transfers both intertemporally and within a period. Starting from an incentive compatible mechanism, how can move transfers across types and across time in a way that still preserves incentives? Lemma 2 answers this question.

The change of variables from $\langle \mathbf{p}, \mathbf{x} \rangle$ to $\langle \mathbf{p}, \mathbf{U} \rangle$ proves key in establishing this result.[32] If we work in an environment with stage transfers it is hard to keep track of the change in incentives caused by a moving transfers around after any given history. But, moving expected utility vectors endogenously keeps the incentives intact. The bijection from $\mathbf{x}$ to $\mathbf{U}$ through $\mathbf{p}$ precisely determines the associated stage transfers.

Now, we construct the collateral dynamic VCG mechanism.

Constructing the Dynamic Collateral VGC mechanism

[32]This approach is popular in (static) contract theory. See, for example Laffont and Martimort [2001]. The key difference is that change with stage transfers in the dynamic environment must be with the expected utility variables, rather than the stage utility ones.

Step 1. Start with the modified VCG mechanism, $\left\langle \mathbf{p}^*, \mathbf{U}^{\mathbf{Mvcg}} \right\rangle$, as defined above. It is incentive compatible and individually rational.

Step 2. Select the mechanism $\langle \mathbf{p}^*, \mathbf{U}^* \rangle$, where $\mathbf{U}^*$ is chosen so that $\inf_{v \in \mathcal{V}} U_B^*(v, c_t | h^{t-1}) = 0 = \inf_{c \in \mathcal{C}} U_S^*(v_t, c | h^{t-1})$ for all v_t, c_t and h^{t-1}. Let $EBS^*(h^{t-1})$ represent the expected budget surplus generated by this mechanism after history h^{t-1}.

Step 3. Show that an incentive compatible and individually rational mechanism guaranteeing efficient trade under interim budget balance can exist if and only if $\langle \mathbf{p}^*, \mathbf{U}^* \rangle$ runs an expected budget surplus, that is, $EBS^*(h^{t-1}) \geq 0 \ \forall h^{t-1}$, $\forall t$.

Step 4. Using $\langle \mathbf{p}^*, \mathbf{U}^* \rangle$ and equation (2.1), recover the stage transfers $\mathbf{x}^*$. $\square$

They key elements of our construction are the minimization of information rents for all types and participation rents to the lowest possible type, and transfer of all remaining surplus to the mechanism designer. We showed in Lemma 1 that local incentive constraints are tight which ensures that the minimalist information rent is paid to the agents. And, using Lemma 2 we extract at every history the excess participation rent generated by the modified VCG mechanism without affecting the incentive constraints. We can now state the main result of the paper.

Proposition 3. *There exists an incentive compatible and individually rational mechanism that implements the efficient allocation under interim budget balance if and only if* $EBS^*(h^{t-1}) \geq 0 \ \forall h^{t-1}$, $\forall t$.

The result is precise in the sense that it offers an "if and only" condition on the primitives of the problem as to when efficiency can be attained. It also signifies the salience of the modified VCG mechanism in characterizing efficient mechanisms in our discrete type dynamic Markovian framework.

Finally we show that the modified VCG mechanism can also be used to characterize efficiency under the traditional ex post budget balance. The key to this result is sustaining efficiency under a static budget surplus every period. Define

$$EBS_t(h^{t-1}) = \mathbb{E}\left[x_B(v_t, c_t|h^{t-1}) - x_{S,t}(v_t, c_t|h^{t-1})|h^{t-1}\right]$$

$$= EBS(h^{t-1}) - \delta\mathbb{E}\left[EBS(h^t)|h^{t-1}\right]$$

Generalizing a standard result on the "equivalence" of ex ante and ex post budget balance in static mechanism design, it is easy to show that there exists a mechanism that implements an allocation under ex post budget balance if and only if it can be implemented by a mechanism that satisfies $EBS_t(h^{t-1}) \geq 0$ for all t.

The modified VCG mechanism extracts all possible rents from the agents. If this mechanism produces a static budget surplus, that is $EBS_t^*(h^{t-1}) \geq 0$ for all h^{t-1}, we are done. If not, than Corollary 3 below shows that as long as $EBS^*(h^{t-1}) \geq 0$ for all h^{t-1} transfers can be moved across time and types using Lemma 2 to produce a mechanism that indeed produces a static budget surplus. The result follows by (backward) induction in a finite horizon model. Infinite horizon model requires a bit more detailed construction.

Corollary 3. *There exists an incentive compatible and individually rational mechanism that implements the efficient allocation under ex post budget balance if and only if* $EBS^*(h^{t-1}) \geq 0$ $\forall h^{t-1}$, $\forall t$.

Unifying existing results, Proposition 5 and Corollary 3 provide a complete characterization of efficiency for dynamic mechanism design models with budget balance. Athey and Miller [2007] provide the characterization under iid shocks. Athey and Segal [2007b, 2013] provide a dynamic version of the AGV mechanism that does not satisfy individual rationality at every history.[33] Finally, Skrzypacz and Toikka [2013]

[33] Athey and Segal [2007a] provide a partial characterization with individual

provide a mechanism that satisfies individual rationality but balances the budget only ex ante.

1.6 Comparative statics

A poignant follow up question to Proposition 5 is of course– what does this condition mean in terms of the parameters of the model? Define

$$\Gamma_t^{Mvcg} = \sum_{\tau=t}^{\infty} \delta^\tau \left(x_{B,\tau}^{Mvcg} - x_{S.\tau}^{Mvcg} \right)$$

Then, $\mathbb{E}\left[\Gamma_t^{Mvcg}|h^{t-1}\right]$ represents the expected cash flow to the mechanism designer from the modified VCG mechanism after history h^{t-1}. From step 2 of the construction of the dynamic collateral VCG mechanism it is easy to see that

$$EBS^*(h^{t-1}) = U_B^{Mvcg}(\underline{v}|h^{t-1}) + U_S^{Mvcg}(\overline{c}|h^{t-1}) + \mathbb{E}\left[\Gamma_t^{Mvcg}|h^{t-1}\right]$$

An intuitive way to think about this definition is to quantify it in the limit continuous type space model. Let

$$\Gamma_t = -\Gamma_t^{vcg} = \sum_{\tau=t}^{\infty} \delta^\tau \left(v_\tau - c_\tau\right)^+$$

where $a^+ = \max\{a, 0\}$. Then, the expected value of Γ_t given the information set at time t (public or private history), defines the expected economic surplus of the dynamic bilateral trade relationship from period t onwards. Let $d \to c$ denote the discrete to continuous limit. In a slight abuse of notation, we have a simple expression for the expected budget surplus in the dynamic collateral VCG mechanism for the

rationality with stronger sufficient conditions. A dynamic version of the pivot mechanism is presented by Bergemann and Valimaki [2010] that satisfies individual rationality every period but not budget balance.

continuous model.

$$\lim_{\mathrm{d} \to \mathrm{c}} EBS^*(h^{t-1}) = \mathbb{E}\left[\Gamma_t | h^{t-1}, v_t = \underline{v}\right] + \mathbb{E}\left[\Gamma_t | h^{t-1}, c_t = \overline{c}\right] - \mathbb{E}\left[\Gamma_t | h^{t-1}\right]$$

Thus, the condition is accounting for the added cash flow equal to the expected surplus generated by the “lowest" types of the buyer and the seller respectively, minus deficit created by the VCG mechanism.

Now, the ability of the underlying environment to sustain efficiency depends on three key factors: extent of overlap of supports of buyer and seller types, level of discounting, persistence of the Markov processes. The necessary and sufficient condition provided by Proposition 5 in the terms of the expected budget surplus of the dynamic collateral VCG mechanism puts joint restrictions on these three aspects of the environment. Of course, given any (or all) of these, relaxing the institutional requirements of individual rationality or budget balance makes efficiency more permissible.

The role of overlap of support is the same as pointed out by Myerson and Satherthwaite [1983] in the static model. If $v_1 > c_M$, then the problem is trivial and trade always happens. On the other hand if $v_N < c_1$, then there is common knowledge of no gains from trade. As the support overlap moves from the former case to the latter, the trade region (weakly) shrinks. Dynamics make the trade regions shift over time but the basic intuition stays.

Fixing the supports, the tension between discounting and persistence drives the potential ability of the dynamic model to break away from the static impossibility results and produce greater levels of efficiency. Two benchmarks are in order. At $\delta = 0$, for any Markov matrix, the model is static and we are back to Myerson and Satterthwaite [1983], and Manea and Kos [2009]. On the other hand, for the Identity Markov matrix, that is for perfectly persistent types, since this a model with commitment, we are back to the repetition of the static optimum for any level of discounting. So, we can converge to the static model at both ends of the parameter space.

Intuitively, it is clear that ceteris paribus, a higher δ is good for efficiency and higher persistence is bad. Higher δ creates more future surplus which can then be used as collateral to sustain efficiency. In the dynamic model, with persistent types, the agents' type not only determines their payoff today but also provides conditional information about payoffs in the future. Thus, as the level of persistence increase so does the amount of asymmetric information about future payoffs. It is thus reasonable to expect that higher persistence make matters wore for efficiency.

Disentangling the two effects in generality, however, is a no mean task. Skrzypacz and Toikka [2013] make some progress on this for the model with ex ante budget balance by imposing certain restrictions on the types and stochastic process. The problem is particularly difficult when looking at interim (or ex post) budget balance because persistence affects both the distortions (information rents) and probability of good or bad shocks every period.

Example 1. We start simple by parameterizing the infinite horizon version of the two type model presented in section 1.4. As before $v_H > c_H > v_L > c_L$. We assume a uniform prior and following Markov matrices.

$$f(v_i|v_i) = g(c_i|c_i) = \alpha, \quad f(v_j|v_i) = g(c_j|c_i) = 1 - \alpha \text{ for } i \neq j$$

So, α measures the level of persistence in the model. The numerical values of types are chosen so that efficiency is not sustainable in the static model, that is, $EBS^*(\delta = 0) < 0$. Since the mechanism is stationary starting period 2 onwards, we need only to calculate five values to test Proposition 3, viz. EBS^*, $EBS^*(v_H, c_H)$, $EBS^*(v_H, c_L)$,

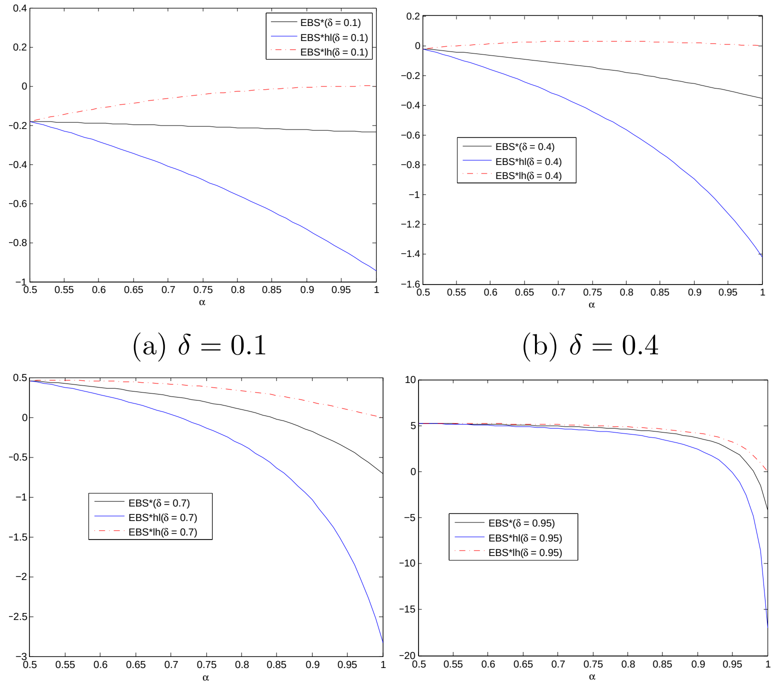

(a) $\delta = 0.1$ (b) $\delta = 0.4$

(c) $\delta = 0.7$ (d) $\delta = 0.95$

Figure 1.2: EBS^*, $EBS^*(v_H, c_L)$ and $EBS^*(v_L, c_H)$ for the two types infinite horizon model.

$EBS^*(v_L, c_H)$, $EBS^*(v_L, c_L)$. For this simple model, these can actually be solved in closed form.

Figure 1.2 plots EBS^*, $EBS^*(v_H, c_L)$ and $EBS^*(v_L, c_H)$ against δ for four different values of δ.

Example 2. Next, we look at the limit continuous type space model where both the buyer and seller types are distributed over $[0, 1]$. The prior is uniform. From the second period onwards types are constant with probability α and are drawn again from the prior with probability $1 - \alpha$. Then, again it is easy to solve the expected budget surplus constraints in closed form. We get

$$EBS^* \geq 0 \Leftrightarrow 1 - \alpha \geq \frac{1-\delta}{\delta}, \text{ and}$$

$$\min_{\{v \in [0,1], c \in [0,1]\}} EBS^*(v, c) \geq 0 \Leftrightarrow \frac{1}{6}\left(\frac{1}{\alpha} - 1\right)^2 \frac{1 + \delta\alpha}{1 - \delta} \geq 1$$

Key economics forces. The intuition that that a higher discount factor is good and higher persistence is bad for efficiency is confirmed by our two examples. Even if the model is arbitrarily close to constant types ($\alpha = 1 - \epsilon$, for ϵ small enough), we can still find a discount factor high enough to achieve efficiency. On the other hand, for any value of δ arbitrarily close to 1, there exists a value of $\alpha < 1$ large enough such that efficiency cannot be attained. So, the limit result on efficiency critically depends of the order of limits.

1.7 Implementation

There is a simple and intuitive method of implementing the dynamic collateral VCG mechanism. Every period both agents deposit a small fees with the mechanism designer post which a VCG style mechanism is run. For all $h^{t-1} \in H^{t-1}$, define

$$r_t^B(h^{t-1}) = U_B(\underline{v}|h^{t-1}) - \delta\mathbb{E}\left[U_B(\underline{v}|h^t) \mid h^{t-1}\right], \text{and}$$

$$r_t^S(h^{t-1}) = U_S(\overline{c}|h^{t-1}) - \delta\mathbb{E}\left[U_S(\overline{c}|h^t) \mid h^{t-1}\right],$$

where the expectation is taken over the t-th element of history $h^t = (h^{t-1}, (v_t, c_t))$ given h^{t-1}. The following simple mechanism implements the dynamic collateral VCG mechanism.

- I At start of every period t (after history h^{t-1}) the agents pay a participation fees $r_t^S(h^{t-1})$ and $r_t^S(h^{t-1})$.
- II Then, run a static modified VCG mechanism.
- III Repeat till time period T.

Recollect that for the two types, two periods model, $r_1^a(h^0)$ and $r_2^a(h^1)$ are exactly the participation fees that we constructed in the simple example in section 1.4.

If the efficient allocation could be implemented by depositing all the extra economic surplus that the bilateral trading relationship hopes to generate right at the start of the contract, then this mechanism would still implement it. However, the converse is not true. If the agents face credit constraints that doesn't allow them to put large sums of money upfront our mechanism can still work.

1.8 Role of an intermediary

If the efficient allocation can be implemented under ex post budget balance, strictly speaking there is no role for an intermediary. Both the agents insure each other against bad shocks; the positive probability of the good shocks and high enough discounting ensure a future economic surplus that can be used as collateral to sustain efficiency. But, as the persistence of bad shocks increases, the size of the transfers along those histories required to sustain efficiency increases rapidly. Thus, with hard bounds on transfers which can interpreted as credit constraints, it is easier to sustain efficiency under interim budget balance rather than ex post.

The role of an intermediary is even more pronounced if we allow for a savings technology, and the interest thus accrued to be carried over in the mechanism designer's budget account, thereby relaxing or strengthening his participation constraint, depending on whether interest is being built over a surplus or a deficit. They key thing to notice is that even in that scenario the collateral dynamic VCG mechanism pins down the necessary and sufficient conditions on efficient implementation. If after history h^{t-1}, the current and future value of the intermediary under the collateral dynamic VCG mechanism is $A^*(h^{t-1})$, then Proposition 5 goes through with a savings technology with $A^*(h^{t-1})$ replacing $EBS^*(h^{t-1})$.

1.9 Conclusion

This paper provides a necessary and sufficient condition for the attainment of efficiency in a repeated bargaining problem with two sided asymmetric information. In the process, it introduces an intermediate notion of budget balance which allows for the role of an intermediary but with bounded credit lines. ced that allows the budget to be balanced dynamically, borrowing from the future but in a bounded fashion. Through a set of simple examples we explore the comparative statics of the underlying economics forces of discounting and level of asymmetric information.

The main result (Proposition 5) and the mechanism can be generalized to ex post notions of incentive compatibility without any cost. In another paper, Lamba [2014], we state and prove a general dynamic payoff equivalence result for the N-player dynamic mechanism design problem with continuous types which is exploited to provide necessary and sufficient conditions for the implementation of the efficient allocation for the general model.

1.10 Appendix

1.10.1 Proof of Lemma 1

We prove the equality for the buyer. The proof for the seller's utility is analogous. Let $u_B^{Mvcg}(v_i)$ be the (expected) current utility of the buyer of type v_i in the modified VCG mechanism. Since the mechanism is stationary we suppress the history. First note that

$$x_B^{Mvcg}(v_i) = \sum_{k=1}^{i} \left[G(v_k) - G(v_{k-1})\right] v_k$$

Then, we have

$$
\begin{aligned}
u_B^{Mvcg}(v_i) &= v_i p^*(v_i|h^{t-1}) - x_B^{Mvcg}(v_i) \\
&= v_i G(v_i|c_{t-1}) - \sum_{k=1}^{i} \left[G(v_k|c_{t-1}) - G(v_{k-1}|c_{t-1})\right] v_k \\
&= (v_i - v_{i-1})\, G(v_{i-1}|c_{t-1}) - \sum_{k=1}^{i-1} \left[G(v_k|c_{t-1}) - G(v_{k-1}|c_{t-1})\right] v_k \\
&= \Delta v_i p^*(v_{i-1}|h^{t-1}) + u_B^{Mvcg}(v_{i-1})
\end{aligned}
$$

Now, to the dynamic mechanism. Fix h^{t-1}. We have

$$
\begin{aligned}
U_B^{Mvcg}(v_i|h^{t-1}) &= u_B^{Mvcg}(v_{i-1}|h^{t-1}) + \delta \sum_{k=1}^{N} f(v_k|v_i) U_B^{Mvcg}(v_k|h^{t-1}, v_i) \\
&= \Delta v_i p^*(v_{i-1}|h^{t-1}) + u_B^{Mvcg}(v_{i-1}|h^{t-1}) + \delta \sum_{k=1}^{N} f(v_k|v_i) U_B^{Mvcg}(v_k|h^{t-1}, v_i) \\
&= \Delta v_i p^*(v_{i-1}|h^{t-1}) + U_B^{Mvcg}(v_{i-1}|h^{t-1}) + \\
&\qquad \delta \sum_{k=1}^{N} \left(f(v_k|v_i) - f(v_k|v_{i-1})\right) U_B^{Mvcg}(v_k|h^{t-1}, v_i) \\
&= \Delta v_i p^*(v_{i-1}|h^{t-1}) + U_B^{Mvcg}(v_{i-1}|h^{t-1}) + \\
&\qquad \delta \sum_{k=1}^{N} \left(f(v_k|v_i) - f(v_k|v_{i-1})\right) U_B^{Mvcg}(v_k|h^{t-1}, v_{i-1})
\end{aligned}
$$

The last equality follows from the fact the mechanism is stationary and as long expectations about future type realizations are the same (which in this case depend only on c_{t-1}) the expected utility vectors are equal. Thus, $U_B^{Mvcg}(v_k|h^{t-1}, v_i) = U_B^{Mvcg}(v_k|h^{t-1}, v_{i-1})$.

1.10.2 Proof of Lemma 2

Suppose $\langle \mathbf{p}, \mathbf{U} \rangle$ is ex post incentive compatible. Fix h^{t-1}. Then, $U_B(v_t|h^{t-1})$ appears in two kinds of incentive compatibility constraints. First,

$$U_B(v_t|h^{t-1}) \geq U_B(v_t'|h^{t-1}) + (v_t - v_t')p(v_t'|h^{t-1})$$

$$+\delta \sum_{i=1}^{N} U_B(v_{t+1,i}|h^{t-1}, v_t') \left(f(v_{t+1,i}|v_t) - f(v_{t+1,i}|v_t')\right)$$

Clearly, addition of the constants $a_B(c_t|h^{t-1})$ to $U(v_t, c_t|h^{t-1})$ for all $v_t \in \mathcal{V}$, and $c_t \in \mathcal{C}$, leads to addition of $a_B(h^{t-1}) = \mathbb{E}\left[a_B(c_t|h^{t-1})\right]$ to $U(v_t|h^{t-1})$ for all $v_t \in \mathcal{V}$, which does not affect any of these constraints.

Next, fix $v_{t-1} = h^{t-1}_{v,t-1}$. Second, we need to consider the constraints,

$$U_B(v_{t-1}'|h^{t-2}) \geq U_B(v_{t-1}|h^{t-2}) + (v_{t-1}' - v_{t-1})p(v_{t-1}|h^{t-2})$$

$$+\delta \sum_{i=1}^{N} U_B(v_{t,i}|h^{t-2}, v_{t-1}) \left(f(v_{t,i}|v_{t-1}') - f(v_{t,i}|v_{t-1})\right)$$

Again, this leads to addition of $a_B(h^{t-1}) = \mathbb{E}\left[a_B(c_t|h^{t-1})\right]$ to $U(v_t|h^{t-1})$ for all $v_t \in \mathcal{V}$ which drops out of the constraint.

Therefore, linear additions of constants as defined in the lemma preserves incentives.

1.10.3 Proof of Proposition 5

Sufficiency is obvious. If $EBS^*(h^{t-1}) \geq 0$, for all $h^{t-1} \in H^{t-1}$, $\forall t$, then the collateral dynamic VCG mechanism satisfies all the necessary properties, and is one such desired mechanism.

Conversely, we will show that the collateral dynamic VCG mechanism produces the highest expected budget surplus at every history of the mechanism. Consider a relaxed problem with only the local "downward" incentive compatibility constraints and the individually rationality constraint of the lowest type. Given first order stochastic dominance and that the efficient allocation is monotonic, if a mechanism implementing the efficient allocation satisfies all the constraints in the relaxed problem, then it is (globally) incentive compatible, and individually rational.[34]

Fixing the allocation to be $\mathbf{p}^*$, we want to choose the mechanism $m = \langle \mathbf{p}^*, \mathbf{U} \rangle$ that produces the highest value(s) of expected budget surplus. For any history h^{t-1} we have

$$EBS(h^{t-1}) = \mathbb{E}^m \left[\sum_{\tau=t}^{T} \delta^{\tau-1} (v_\tau - c_\tau) p_\tau - U_B(v_t|h^{t-1}) - U_S(c_t|h^{t-1}) \mid h^{t-1} \right]$$

It is thus straightforward to see that in order to produce the highest value of $EBS^*(h^{t-1})$ all the local "downward" incentive constraints and the individual rationality constraint of the lowest type must hold as equalities, which is isomorphic to the collateral dynamic VCG mechanism. Thus, for any incentive compatible and individually rational mechanism $m = \langle \mathbf{p}^*, \mathbf{U} \rangle$, we have

$$U_B(v_t, c_t|h^{t-1}) \geq U_B^*(v_t, c_t|h^{t-1}) \quad \text{and} \quad U_S(v_t, c_t|h^{t-1}) \geq U_S^*(v_t, c_t|h^{t-1})$$

Thus, if for any history $h^{t-1} \in \mathcal{H}$, $EBS^*(h^{t-1}) < 0$, we must have $EBS(h^{t-1}) < 0$ in $\langle \mathbf{p}^*, \mathbf{U} \rangle$. The result follows.

1.10.4 Proof of Corollary 3

First, define,

$$EBS_t(h^{t-1}) = EBS(h^{t-1}) - \mathbb{E}^m \left[EBS(h^t)|h^{t-1} \right] =$$

[34] See Pavan, Segal and Toikka [2014] and Battaglini and Lamba [2014].

$$\mathbb{E}^m\left[x_B(v_t, c_t|h^{t-1}) - x_S(v_t, c_t|h^{t-1})|h^{t-1}\right]$$

to be the current expected budget surplus. We show that implementation under ex post budget balance balance is equivalent to generating a current expected budget surplus at every history.

Lemma 3. *An allocation* $\mathbf{p}$ *is implementable under ex post budget balance if and only if there exists an implementable mechanism* $m = \langle \mathbf{p}, \mathbf{x} \rangle$ *that satisfies* $EBS_t(h^{t-1}) \geq 0$, $\forall h^{t-1}$, $\forall t$.

Proof. Ex post budget balance implies $EBS_t(h^{t-1}) \geq 0$, $\forall h^{t-1}$, $\forall t$ is obvious. Conversely, suppose $\langle \mathbf{p}, \mathbf{x} \rangle$ satisfies $EBS_t(h^{t-1}) \geq 0$, $\forall h^{t-1}$, $\forall t$. Fix a history h^{t-1}, and let $\Pi = EBS_t(h^{t-1}) \geq 0$. Define

$$\tilde{x}(v_t, c_t|h^{t-1}) = x_B(v_t|h^{t-1}) - \sum_{i=1}^{N} x_B(v_{t,i}, c_t) f(v_{t,i}|v_{t-1}) dv_t + x_S(c_t|h^{t-1}) + \alpha\Pi,$$

where $\alpha \in [0, 1]$ is a constant. We have

$$\tilde{x}(v_t|h^{t-1}) = x_B(v_t|h^{t-1}) - (1 - \alpha)\Pi, \text{ and}$$

$$\tilde{x}(c_t|h^{t-1}) = x_S(c_t|h^{t-1}) + \alpha\Pi$$

Repeat this for every possible history. Now, consider the mechanism $\langle \mathbf{p}, \tilde{\mathbf{x}} \rangle$. It is ex post budget balanced by construction. Moreover, using an incentive compatible mechanism we are reducing what the buyer has to pay and increasing what the seller gets. So, the new mechanism must also be incentive compatible and individually rational. □

Now, to the main result. Note that If there exists t and h^{t-1} such that $EBS^*(h^{t-1}) < 0$, then by Proposition 5, the efficient allocation cannot be implemented under interim budget balance, and hence neither under ex post budget balance.

Conversely, suppose $EBS^*(h^{t-1}) \geq 0$ $\forall h^{t-1}$, $\forall t$. We show that then, there exists a mechanism $\left\langle \mathbf{p}, \tilde{\mathbf{U}} \right\rangle$ that satisfies $EBS_t(h^{t-1}) \geq 0$ $\forall h^{t-1}$. If T is finite the result follows from backward induction. So, let $T = \infty$.

Define

$$\mathcal{H}^1 = \left\{ h^{t-1} \in \mathcal{P}\left(H\right) | EBS^*_\tau(h^{\tau-1}) \geq 0 \; \forall \; \tau \leq t, h^{\tau-1} = \pi_{\tau-1}(h^{t-1}) \right\}$$

and,

$$\mathcal{H}^2 = \left\{ h^{t-1} \in \mathcal{P}\left(H\right) | EBS^*_t(h^{t-1}) + \mathbb{E}\Big[\sum_{\tau=t+1}^{T} \mathbf{1}_{\{h^{t-1} \in \mathcal{H}^1\}} \delta^{\tau-t} EBS^*_\tau(h^{\tau-1}) \mid h^{t-1} \Big] \geq 0 \right\}$$

$$\mathcal{H}^3 = \left\{ h^{t-1} \in \mathcal{P}\left(H\right) | EBS^*_t(h^{t-1}) + \mathbb{E}\Big[\sum_{\tau=t+1}^{T} \mathbf{1}_{\{h^{t-1} \in \mathcal{H}^2\}} \delta^{\tau-t} EBS^*_\tau(h^{\tau-1}) \mid h^{t-1} \Big] \geq 0 \right\}$$

.

.

.

$$\mathcal{H}^n = \left\{ h^{t-1} \in \mathcal{P}\left(H\right) | EBS^*_t(h^{t-1}) + \mathbb{E}\Big[\sum_{\tau=t+1}^{T} \mathbf{1}_{\{h^{t-1} \in \mathcal{H}^{n-1}\}} \delta^{\tau-t} EBS^*_\tau(h^{\tau-1}) \mid h^{t-1} \Big] \geq 0 \right\}$$

Then, is clear that $\lim_{n \to \infty} \mathcal{H}^n = \mathcal{P}(H)$. Finally, construct the following mechanism.

Algorithm for a mechanism $\langle \mathbf{p}^*, \bar{\mathbf{U}} \rangle$ that satisfies ex post budget balance

Step 1. Start with the collateral dynamic VCG mechanism. If $EBS^*_t(h^{t-1}) \geq 0$ $\forall h^{t-1}$, $\forall t$ we are done. If not go to step 2.

Step 2.1. Fix a history h^{t-1} such that $EBS^*(h^{t-1}) < 0$. Then, for all histories $h^{\tau-1} \in \mathcal{H}^1$ where $\tau > t+1$ and $\pi_{t-1}(h^{\tau-1}) = h^{t-1}$, increase $U_B(v_\tau, c_\tau | h^{\tau-1})$ and

$U_S(v_\tau, c_\tau|h^{\tau-1})$ so that $EBS(h^{\tau-1}) = 0$. All the surplus backloaded is then utilized by increasing $U_B(v_{t+1}, c_{t+1}|h^t)$ and $U_S(v_{t+1}, c_{t+1}|h^t)$ using Lemma 2. Do this for all histories for which $EBS^*(h^{t-1}) < 0$. If for those histories $EBS^*(h^{t-1}) \geq 0$, we are done. If not, go to step 2.2....

Step 2.n. Do, the same as above but change is made for all histories $h^{\tau-1} \in \mathcal{H}^n$ where $\tau > t+1$ and $\pi_{t-1}(h^{\tau-1}) = h^{t-1}$.

Step 3. Since $EBS^*(h^{t-1}) \geq 0$ $\forall h^{t-1}$, $\forall t$, the sequences of mechanisms constructed in step 2 converges to a finite mechanism that satisfies $EBS_t^*(h^{t-1}) \geq 0$ $\forall h^{t-1}$, $\forall t$.

Step 4. Using Lemma 3 construct a mechanism that satisfies ex post budget balance. □

Bibliography

[1] Claude d'Aspremont and Louis-André Gérard-Varet (1979a). "Incentives and Incomplete Information," *Journal of Public Economics*, 11(1): 25-45.

[2] Susan Athey and David Miller (2007). "EfÞciency in Repeated Trade with Hidden Valuations," *Theoretical Economics*, 2(3): 299-354.

[3] Susan Athey and Ilya Segal (2007a). "Designing Efficient Mechanisms for Dynamic Bilateral Trading Games," *American Economic Review P&P* , 97(2): 131-136.

[4] Susan Athey and Ilya Segal (2007b). "An Efficient Dynamic Mechanism," *working paper.*

[5] Susan Athey and Ilya Segal (2013). "An Efficient Dynamic Mechanism," *Econometirca, forthcoming.*

[6] Marco Battaglini (2005). "Long Term Contracting with Markovian Consumers," *American Economic Review*, 95(3): 637-658.

[7] Marco Battaglini and Rohit Lamba (2014). "Optimal Dynamic Contracting: the First-Order Approach and Beyond," *working paper.*

[8] Dirk Bergemann and Maher Said (2011). "Dynamic Auctions: A Survey," *Wiley Encyclopedia of Operations Research and Management Science.*

[9] Dirk Bergemann and Juuso Valimaki (2010). "The Dynamic Pivot Mechanism," *Econometrica*, 78(2): 771-789.

[10] Kalyan Chatterjee and William Samuelson (1983). "Bargaining under Incomplete Information," *Operations Research*, 31, 835-851.

[11] Ronald Coase (1960). "The Problem of Social Cost," *Journal of Law and Economics*, 3: 1-44.

[12] Theodore Groves (1973). "Incentives in Teams," *Econometrica*, 41(4): 617-631.

[13] Mathew Jackson (2003). "Mechanism Theory," *working paper.*

[14] Vijay Krishna and Motty Perry (2000). "Efficient Mechanism Design,"*working paper.*

[15] Rohit Lamba (2014). "Dynamic Payoff Equivalence and Efficient Mechanism Design," *working paper.*

[16] Jonathan Levin (2003). "Relational Incentive Contracts," *American Economic Review*, 93(3): 835-357.

[17] Lars Ljungqvist and Thomas Sargent (2004). *Recursive Macroeconomic Theory,* MIT Press.

[18] Mailath, G. and A. Postelwaite (1991). "Asymmetric Information Bargaining Problems with Many Agents," *Review of Economic Studies*, 57(3), 351-367.

[19] Andreu Mas-Colell, Michael Whinston, and Jerry Green (1995). *Microeconomic Theory,* Oxford University Press.

[20] Toshihide Matsuo (1989). "On Incentive Compatible, Individually Rational, and Ex Post Efficient Mechanisms for Bilateral Trading." *Journal of Economic Theory*, 29(2): 265-281.

[21] Roger Myerson (1985). "Analysis of Two Bargaining Problems with Incomplete Information," *Game Theoretic Models of Bargaining*, edited by Al Roth, Cambridge University Press.

[22] Alessandro Pavan, Ilya Segal and Juuso Toikka (2014). "Dynamic Mechanism Design: A Myersonian Approach," *Econometirca,* 82(2): 601-653.

[23] Jonathan Thomas and Tim Worrall (1988). "Self-Enforcing Wage Contracts," *Review of Economic Studies*, 55(4): 541-553.

[24] William Vickrey (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16(1): 8-37.

[25] Rakesh Vohra (2012). "Dynamic Mechanism Design," *Surveys in Operations Research and Management Science*, 17: 60-68.

[26] Steven R. Williams (1999). "A Characterization of Efficient, Bayesian Incentive Compatible Mechanisms," *Economic Theory*, 14: 155-180.

Chapter 2

Dynamic Payoff Equivalence and Efficient Mechanism Design

2.1 Introduction

Questions of design in real economic situations are often dynamic. Information and actions are reported and recorded. Managerial compensation and taxation are good examples. Repeated interactions are also being internalized in the design of spectrum auctions. All these demand the economic theory of mechanism design to be adept to changing underlying environments and evolving information. Adjusting existing static results to the dynamic models and introducing new ones is thus what the doctor orders.

A key result in static mechanism design with a strong theoretical appeal and wide set of applications is the *payoff equivalence result*. It essentially states that for a model with quasilinear preferences and independent types if two mechanisms implement the same allocation, then the expected utility (and thus transfers) for each agent in one must differ from the the other through a constant. The result is powerful for it classifies all implementable mechanisms into simple equivalence classes. Often the revenue maximizing or least costly mechanism amongst the class is then chosen to be the one that gives exactly the reservation utility to the “lowest" type, inductively constructing payments for all other types.

It has been an open question in the literature on dynamic mechanism design as to how transfers can be moved across time and across types while preserving incentives.[1] In other words, a history dependent version of the payoff equivalence result has been elusive. The main task of this paper is to fill that gap.

The *dynamic payoff equivalence result* provides a precise characterization of equivalence classes of implementable mechanisms in models where the agents interact more than once. The problem is challenging because tinkering with transfers in period t will affect incentives in periods $s \leq t$ through stochastic evolution of information and the prescribed mechanism. We show that given two mechanisms that implement the same allocation, expected utility of an agent after any history in one must differ from the other through a history dependent constant.

The analysis is simplified by exploiting a dynamic version of the revelation principle. It allows us to, without loss of generality, look at mechanisms where the agents report their type to the mechanism every period. Since private information of the agents may arrive (and change) over time, a moral call on what the mechanism itself reveals to the agents needs to be taken. We look at the two extreme points of the information set: *public mechanisms* where the all agents can see each others' announcements and *blind mechanisms* where the agent does not observe the other players' announcements.

Armed with the dynamic payoff equivalence result, we explore the implementation of (ex post) efficient allocation in dynamic mechanism models. A series of papers have

[1]For instance, constructing a partial characterization of efficiency, Athey and Segal [2007a] write “The only degree of freedom the transfers offer in transferring utility across players is a fixed constant K (if it varied with history, it would affect incentives)."

explored the implementability of the efficient allocation under various institutional (participation, budget balance, efficient exit) and informational (Bayesian and ex post incentive compatibility) constraints. Since the allocation is fixed to be efficient, our result helps unify these existing results simply through the family of history dependent constants. In particular it helps us recast well known impossibility results on implementation of the efficient allocation under budget balance and participation constraints into precise *(im)possibility* characterizations for the dynamic models.[2]

Starting with a dynamic version of VCG mechanism, which does not satisfy budget balance, we construct a mechanism for a general N-player mechanism design problem which precisely characterizes the implementability of the efficient allocation under budget balance. Finally, we also use the dynamic payoff equivalence result to characterize the second-best mechanism for the dynamic bilateral trading problem. For a simple two period iid model, we also precisely characterize the no trade regions. This represents but a small set of a wide array of applications where the dynamic payoff equivalence formula can be put to use to calculate the revenue or social welfare maximizing contracts.

Related Literature. The paper builds on a body of work from static mechanism design. Exploring optimal mechanisms, Myerson [1981], Wilson [1993] and others provide revenue and payoff equivalence results. Vickery [1961] and Groves [1973] were the early papers to talk about efficient mechanisms. We refer the reader to Mas-Collel, Whinston and Green [1995] and Milgrom [2004] for detailed overviews.

Bergemann and Valimaki [2010], Pavan, Segal and Toikka [2014], and Skrzypacz and Toikka [2013] have presented dynamic versions of the payoff equivalence result in different contexts and notions of incentive compatibility. Our result is both more general and nuanced. First it encapsulates all these existing results as special cases.

[2]See Myerson and Satterthwaite [1983] and Mailath and Postlewaite [1991] for impossibility results in private and and public goods environments respectively.

Second, and more importantly, versions of the dynamic payoff equivalence result presented before this were ex ante ones that pin down the ex ante aggregate value of transfers through the entire length of time up to a constant. But, they are silent on how can these transfers vary over time. For example, these results cannot tell you how expected utility is related from period 2 onwards in two mechanisms that implement the same allocation. Thus, what is the appropriate equivalence class of mechanisms that one can consider in the dynamic framework to characterize efficient and revenue maximizing contracts? Our dynamic and history dependent version of the payoff equivalence result answers these questions precisely.

Athey and Segal [2007b, 2013] and Lamba [2014] provide mechanisms that implement the efficient allocation under budget balance. The former generalizes the AGV mechanism to the dynamic environment and thus does not satisfy individual rationality. And restricting itself to the bilateral trade setting, the latter provides precise conditions on the primitives of the model under which efficiency can be attained under both budget balance and individual rationality.[3] In section 6, we generalize Lamba's results to an N player mechanism design problem using dynamic payoff equivalence.

2.2 Model

There are N agents who commit to participate in an economic relationship for T discrete periods, where $T \leq \infty$. At the inception of every period, each agent privately observes a payoff relevant information parametrized by $\theta_i \in \Theta_i$ for $i \in \{1, 2, ..., N\}$. We will write $\theta_t = (\theta_{i,t})_{i=1}^N$. The economic relationship is governed by an allocation $k_t \in K$ which is observed by all, and transfers $x_{i,t}$ every period. For an allocation and

[3]Skrzypacz and Toikka [2013] also provide a characterization of efficiency under budget balance and individual rationality. However, they only require the budget to balance ex ante which can lead to the mechanism designer subsidizing the agents by arbitrarily large amounts when the time horizon is long enough.

transfer scheme $(k_t, x_{i,t})_{t=1}^{T}$, preferences are quasilinear and time separable, given by

$$\sum_{t=1}^{T} \delta^{t-1} \left(u_i(\theta_t, k_t) - x_{i,t} \right)$$

where $\delta \in (0,1)$ is the common discount factor, and u_i is stage utility function assumed to bounded,

Types are first independently drawn from some priors $(F_i)_{i=1}^{N}$, and hence from independent Markov processes $(F_i(.|.,.))_{i=1}^{N}$, where $F_i : \Theta_i \times K \mapsto \Theta_i$. For simplicity, we assume that prior and the Markov processes have full support.[4] The expected stage payoff must be bounded– there exists an $M < \infty$ such that for any i, k and θ_t, and $F = \prod_{i=1}^{N} F_i$,

$$\int_{\Theta_t} |u_i(\theta_{t+1}, k)| dF(\theta_{t+1}|\theta_t, k) < M$$

The model encapsulates as a special case the iid model, the perfectly persistent case and the AR(1) model. It also lets the stochastic evolution of types depend on the allocation rule. However, in the rest of the paper we'll look at the special case where the Markov evolution is exogenous, that is, $F_i : \Theta_i \mapsto \Theta_i$. At the cost of lengthy notation, the main results easily extend to the case where stochastic evolution can also depend on past allocations. Finally, the independence assumption across players is made for reasons analogous to the static model: to avoid full surplus extraction a'la Cremer and McLean [1988].

[4]Note that we can include moving supports $\Theta_{i,t}$, time varying Markov processes $F_{i,t}$, and make the Markov process at time t depend on the entire history of allocation K^{t-1}. The chosen setup is simple in communicating the main result and can be easily generalized.

2.3 Mechanisms

It is easy to show that a form of revelation principle holds and thus we can, without loss of generality, consider direct mechanisms. Every period the agents learn their own types, and then send a report to the mechanism, which in turn, spits out the allocation and transfers rules. Employing the revelation principle, however, demands a moral call on the information the mechanism itself reveals to the agents. What is the degree of transparency in the mechanism? In particular, does one agent observe the other agents' announcements?

Most of the literature so far has been silent on the applicability and generality of the revelation principle with respect to information sharing in dynamic mechanism design models.[5] Myerson [1986] has argued that in a general multistage game with communication with private information and publicly observable actions, confidentiality is essential to generate the most permissible results. This is intuitive: private announcements means less leakage of information to the agents and thus less incentive constraints to keep track of.

We shall mostly work in the two extreme environments- one in which all announcements are publicly observed, and the other where there is no release of information beyond the allocation rule, viz. agents cannot see each others announcements and transfers are measurable with respect to their information sets. There is a close information theoretic relationship between these *public* and *blind* mechanisms respectively. Incentives, participation and budget balance constraints satisfied by the public mechanism are of course satisfied by the blind mechanism too, whereas the converse may not always hold.

The set of feasible histories in period t for the mechanism m is given by $H^{m,t}$. In

[5] Skrzypacz and Toikka [2013] is an exception.

general, a typical element of $H^{m,t}$, say h^t, can be recursively defined as follows.

$$h^0 = \emptyset, \quad h^t = \left\{h^{t-1}, (\hat{\theta}_{i,t})_{i=1}^N, k_t\right\}$$

where $\hat{\theta}_{i,t}$ is the announcement by agent i at time t. Similarly, the set of possible histories for the agents, $H_i^{m,t}$ can be defined as follows. For the public mechanism,

$$h_i^{pub,1} = \{\theta_{i,1}\}, \quad h_i^{pub,t} = \left\{h_i^{pub,t-1}, (\hat{\theta}_{k,t-1})_{k=1}^N, \theta_{i,t}, k_t\right\},$$

where $\theta_{i,t}$ is the actual type of agent i at time t. And, for the blind mechanism,

$$h_i^{blind,1} = \{\theta_{i,1}\}, \quad h_i^{blind,t} = \left\{h_i^{blind,t-1}, \hat{\theta}_{i,t-1}, \theta_{i,t}, k_t\right\}$$

As we will see below, observing all the announcements leads to knowing the allocation in the direct mechanism, so writing down k_t in the agent's information set for the public mechanism is superfluous. However, the same cannot be said for the blind mechanism, where the allocation rule can carry additional information.

In what follows, we suppress the specification of the mechanism when it is obvious. Also, under truthful histories the difference between announcements and actual types will be moot and thus suppressed.

The direct mechanism, say m, is then a collection of history dependent allocation and transfer vectors, $m = \langle \mathbf{k}, \mathbf{x} \rangle = \left(k\left(h^t\right), x_i\left(h^t\right)_{i=1}^N\right)_{t=1}^T$. To differentiate current and past types, we will often use the alternative notation $k(\theta_t|h^{t-1})$. Moreover, for a given mechanism, the strategy for agent i, $(\sigma_i^{m,t})_{t=1}^T$, is simply a function that maps private history into an announcement every period- $\sigma_i^{m,t} : H_i^{m,t} \mapsto \Theta_i$.

In most of what follows, we present further notations and definitions in terms of the public mechanism. This is mostly done for two reasons. First, for the ease of notation– for a public mechanism under truthful histories, the history of the mechanism coincides with the private history of the agents. Second, though a realistic level of transparency for a mechanism perhaps lies in the middle of the two extreme cases being presented, the blind mechanism seems to be a highly artificial construct, interesting only as a benchmark. The definitions and results for the blind mechanism are presented in the appendix.

Additional Notation. Along truthful histories, we can write the stage utility succinctly as

$$u_i(\theta_t, k(\theta_t|h^{t-1})) = u_i(\theta_t|h^{t-1})$$

In order to keep notation simple we suppress the type/variable over which expectation is taken. For example

$$u_i(\theta_{i,t}|h^{t-1}) = \mathbb{E}_{\theta_{-i,t}}\left[u_i\left(\theta_{i,t}, \theta_{-i,t}\left|h^{t-1}\right.\right)\left|h^{t-1}\right.\right]$$

$$u_i(\theta_{i,t+1}|h^{t-1}, \theta_{i,t}) = \mathbb{E}_{\theta_{-i,t+1}}\left[\mathbb{E}_{\theta_{-i,t}}\left[u_i(\theta_{i,t+1}, \theta_{-i,t+1}\left|h^{t-1}, \theta_{i,t}, \theta_{-i,t}\right.)\left|h^{t-1}, \theta_{i,t}\right.\right]\right]$$

Expected utility of agent i can be recursively defined as

$$U_i(\theta_t|h^{t-1}) = u_i(\theta_t|h^{t-1}) - x_i(\theta_t|h^{t-1}) + \delta\mathbb{E}_{\theta_{t+1}}\left[U_i(\theta_{t+1}|h^{t-1}, \theta_t)\left|h^{t-1}, \theta_t\right.\right] \qquad (2.1)$$

and,

$$U_i(\theta_{i,t}|h^{t-1}) = u_i(\theta_{i,t}|h^{t-1}) - x_i(\theta_{i,t}|h^{t-1}) + \delta\mathbb{E}_{\theta_{t+1}}\left[U_i(\theta_{t+1}|h^{t-1}, \theta_{i,t})\left|h^{t-1}, \theta_{i,t}\right.\right]$$

Let $u_i(\theta'_{i,t}; \theta_t|h^{t-1}) = u_i(\theta_t, k(\theta'_{i,t}, \theta_{-i,t}|h^{t-1}))$. Then, utility of agent i of type $\theta_{i,t}$ from misreporting (once) to be type $\theta'_{i,t}$, for a fixed type $\theta_{-i,t}$ of the other agents, can be succinctly written as

$$U_i(\theta'_{i,t};\theta_t|h^{t-1}) = u_i(\theta'_{i,t};\theta_t|h^{t-1}) - x_i(\theta'_{i,t},\theta_{-i,t}|h^{t-1})$$

$$+\delta\mathbb{E}_{\theta_{t+1}}\left[U_i(\theta_{t+1}|h^{t-1},\theta'_{i,t},\theta_{-i,t})\,\middle|h^{t-1},\theta_t\right]$$

$$= U_i(\theta'_{i,t},\theta_{-i,t}|h^{t-1}) + \big(u_i(\theta'_{i,t};\theta_t|h^{t-1}) - u_i(\theta'_{i,t},\theta_{-i,t}|h^{t-1})\big)$$

$$+\delta\int\limits_{\Theta_i} U_i(\theta_{i,t+1}|h^{t-1},\theta'_{i,t},\theta_{-i,t})\big(dF_i(\theta_{i,t+1}|\theta_{i,t}) - dF_i(\theta_{i,t+1}|\theta'_{i,t})\big) \tag{2.2}$$

Similarly,

$$U_i(\theta'_{i,t};\theta_{i,t}|h^{t-1}) = U_i(\theta'_{i,t}|h^{t-1}) + \big(u_i(\theta'_{i,t};\theta_{i,t}|h^{t-1}) - u_i(\theta'_{i,t}|h^{t-1})\big)$$

$$+\delta\int\limits_{\Theta_i} U_i(\theta_{i,t+1}|h^{t-1},\theta'_{i,t})\big(dF_i(\theta_{i,t+1}|\theta_{i,t}) - dF_i(\theta_{i,t+1}|\theta'_{i,t})\big)$$

Mechanism. *It is straightforward to note that a mechanism $m = \langle \mathbf{k}, \mathbf{x}\rangle$, which is a collection of history dependent allocation and transfer vectors, can be equivalently defined to be $m = \langle \mathbf{k}, \mathbf{U}\rangle$, a collection of history dependent allocation and expected utility vectors, where (fixing the allocation) the duality between transfers and expected utility vectors is completely described by equation (2.1).*

Incentive Compatibility. Exploiting the one-deviation principle, incentive compatibility can be defined as follows.[6]

Definition 7. *A mechanism $m = \langle \mathbf{k}, \mathbf{U}\rangle$ satisfies perfect Bayesian incentive compat-*

[6]See Pavan, Segal and Toikka [2014] for the validity of the one-deviation principle here.

ibility if truthtellng is optimal at all truthful histories, that is,

$$U_i(\theta_{i,t}|h^{t-1}) \geq U_i(\theta'_{i,t};\theta_{i,t}|h^{t-1})$$

$\forall i \in \{1,...,N\}$, $\forall \theta_{i,t}, \theta'_{i,t} \in \Theta_i$, $\forall h^{t-1} \in H^{t-1}$, $\forall t$.

In addition, in the perfect Bayesian equilibrium of the associated game, every agent believes that the other agents are following a truthful strategy. Moreover, the Markov assumption means that it is incentive compatible for an agent to say the truth even if she/he has lied in the past.

A *stronger* equilibrium notion is the that of ex-post incentive compatibility. The mechanism in each period is implemented in ex-post equilibrium (see Chung and Ely [2006]). Formally,

Definition 8. *A mechanism $m = \langle \mathbf{k}, \mathbf{U} \rangle$ satisfies ex post incentive compatibility if*

$$U_i(\theta_t|h^{t-1}) \geq U_i(\theta'_{i,t};\theta_t|h^{t-1})$$

$\forall i \in \{1,...,N\}$, $\forall \theta_{i,t}, \theta'_{i,t} \in \Theta_i$, $\forall \theta_{-i,t} \in \Theta_{-i}$, $\forall h^{t-1} \in H^{t-1}$, $\forall t$.

In the words of Bergemann and Valimaki [2010], "We say that the dynamic direct mechanism is periodic ex post incentive compatible if truthtelling is a best response regardless of the history and the current state of the other agents."

2.4 Dynamic Payoff Equivalence

Now to the dynamic and history dependent version of the payoff equivalence result. To fix ideas, consider a two period model. Suppose $m = \langle \mathbf{k}, \mathbf{x} \rangle$ is an incentive compatible mechanism. Change the transfers to all possible types of agent i in period 2 after history θ_1 by a constant say a_i. Standing at history θ_1, (as in the static model) this change in the mechanism does not affect incentives in period 2. However, the same cannot be said for period 1. The change renders the contract not incentive compatible in period 1– it changes the value of $U_i(\theta_{i,1})$ keeping utility of all other agent i types the same. The transfers in period 1 for type $\theta_{i,1}$ thus need to be adjusted to keep the contract incentive compatible.

Think of a general T period model. Keeping track of these history dependent transfers to pin down the equivalence class of implementable allocations can be a daunting task. A change of variable though makes the problem simpler. We look at the mechanism through the prism of allocation and expected utility vectors, $\langle \mathbf{k}, \mathbf{U} \rangle$.

A dynamic analog of the payoff equivalence result then follows. The set of incentive compatible allocations sits neatly in a family of history dependent constants. Starting from any incentive compatible mechanism, scaling expected utility vectors through constants that may depend on the history and current type of other agents, but not the current type of the given agent produces another incentive compatible mechanism. Moreover, all incentive compatible mechanisms must belong to this class of scaled family of history dependent constants.

Proposition 4. *Payoff equivalence holds after every history. That is, if $\langle \mathbf{k}, \mathbf{U} \rangle$ and $\left\langle \mathbf{k}, \tilde{\mathbf{U}} \right\rangle$ are two ex post incentive compatible mechanisms that generate utility vectors $\left((U_i(\theta_t|h^{t-1}))_{i=1}^N \right)_{t=1}^T$ and $\left(\left(\tilde{U}_i(\theta_t|h^{t-1}) \right)_{i=1}^N \right)_{t=1}^T$ respectively, then, there exists a family of constants $\left((a_i(\theta_{-i,t}|h^{t-1}))_{i=1}^N \right)_{t=1}^T$ such that*

$$U_i(\theta_t|h^{t-1}) = \tilde{U}_i(\theta_t|h^{t-1}) + a_i(\theta_{-i,t}|h^{t-1})$$

Conversely, if $\langle \mathbf{k}, \mathbf{U} \rangle$ is ex post incentive compatible, and $\mathbf{U}$ and $\tilde{\mathbf{U}}$ satisfy the above two equations for a finite family of constants $\left((a_i(\theta_{-i,t}|h^{t-1}))_{i=1}^N \right)_{t=1}^T$, then $\left\langle \mathbf{k}, \tilde{\mathbf{U}} \right\rangle$ is also ex post incentive compatible.

A natural next question question to ask is: what does the result mean in terms of actual transfers? Let us go back to the two period model. Suppose as before that we increase the transfers to all possible types of agent i in period 2 after history θ_1 by a constant say a_i. This translates into an increase in $U_i(\theta_2|\theta_1)$ by a_i for $\theta_2 \in \Theta$. Rewrite equation (2.1) for this two period model:

$$U_i(\theta_1) = u(\theta_1) - \tilde{x}_i(\theta_1) + \delta \mathbb{E}_{\theta_2} \left[U_i(\theta_2|\theta_1) + a_i(\theta_1) \, |\theta_1 \right]$$

Keeping $U_i(\theta_1)$ the same as before, this uniquely defines the the new transfer $\tilde{x}_i(\theta_1)$that keeps the mechanism incentive compatible. This change was of course a very simple one, done for only one history and one agent and independent of the other agent types. The more detailed the required change, the more equation (2.1) will be put to use to derive transfers for the “new" mechanism.

Note that the above result was established for a fairly strong notion of incentive compatibility. As a simple corollary, one can show that an analogous payoff equivalence result also holds for perfect Bayesian incentive compatibility.

Corollary 4. *If $\langle \mathbf{k}, \mathbf{U} \rangle$ and $\left\langle \mathbf{k}, \tilde{\mathbf{U}} \right\rangle$ are two perfect Bayesian incentive compatible mechanisms that generate utility vectors $\left((U_i(\theta_t|h^{t-1}))_{i=1}^N \right)_{t=1}^T$ and $\left(\left(\tilde{U}_i(\theta_t|h^{t-1}) \right)_{i=1}^N \right)_{t=1}^T$ respectively, then, there exists a family of constants $\left((a_i(h^{t-1}))_{i=1}^N \right)_{t=1}^T$ such that*

$$U_i(\theta_{i,t}|h^{t-1}) = \tilde{U}_i(\theta_{i,t}|h^{t-1}) + a_i(h^{t-1})$$

Conversely, if $\langle \mathbf{k}, \mathbf{U} \rangle$ is perfect Bayesian incentive compatible, and $\mathbf{U}$ and $\tilde{\mathbf{U}}$ satisfy the above two equations for a finite family of constants $\left((a_i(h^{t-1}))_{i=1}^N \right)_{t=1}^T$, then $\left\langle \mathbf{k}, \tilde{\mathbf{U}} \right\rangle$ is also perfect Bayesian incentive compatible.

All versions of the dynamic payoff equivalence result presented before this were ex ante ones that pin down the ex ante aggregate value of transfers through period 1 to T

up to a constant.[7] But, they are silent on how can these transfers vary over time. For example, these results cannot tell you how expected utility is related from period 2 onwards in two mechanisms that implement the same allocation. Our dynamic payoff equivalence result answers these questions and identifies the appropriate equivalence class of incentive compatible mechanisms in the dynamic framework that can then be used to characterize efficient and revenue maximizing contracts.

2.5 Other Institutional Constraints

While the dynamic payoff equivalence result characterizes all incentive compatible mechanisms, there are other institutional constraints that capture important aspects of real economic situations in mechanism design. We discuss two commonly used: individual rationality and budget balance.[8]

Individual Rationality. Each agent is allowed to walk away at any stage of the mechanism after learning her/his type that period if the utility from continuing in the contract falls below the reservation threshold, which is normalized to zero.

Definition 9. *A mechanism* $m = \langle \mathbf{k}, \mathbf{U} \rangle$ *satisfies perfect Bayesian individually rationality if*

$$U_i(\theta_{i,t}|h^{t-1}) \geq 0$$

$\forall i \in \{1, ..., N\}$, $\forall \theta_{i,t}, \in \Theta_i$, $\forall h^{t-1} \in H^{t-1}$, $\forall t$.

[7] See Bergemann and Valimaki [2010], Pavan, Segal and Toikka [2014], and Skrzypacz and Toikka [2013].

[8] Athey and Miller [2007] write, "At the outset, we should point out that IC [incentive compatibility] imposes restrictions on equilibria in the game, while the BB [budget balance] and IR [individual rationality] assumptions are better thought of as conditions on the structure of the game itself."

Definition 10. *A mechanism* $m = \langle \mathbf{k}, \mathbf{U} \rangle$ *satisfies ex post individually rationality if*

$$U_i(\theta_t|h^{t-1}) \geq 0$$

$\forall i \in \{1, ..., N\}$, $\forall \theta_t, \in \Theta$, $\forall h^{t-1} \in H^{t-1}$, $\forall t$.

Budget Balance. In mechanism design with many agents budget balance is seen as the limits on insurance or external subsidies available to them. The most widely used notion is that of ex post budget balance.

Definition 11. *A mechanism* $m = \langle \mathbf{k}, \mathbf{x} \rangle$ *satisfies ex post budget balance if*

$$\sum_{i=1}^{N} x_i(\theta_t|h^{t-1}) = 0,$$

$\forall \theta_t \in \Theta$, $\forall h^{t-1} \in H^{t-1}$, $\forall t$.

Lamba [2014] introduce an intermediate notion that allows for the budget to balanced dynamically, not running a deficit in expectation after any history.

Definition 12. *A mechanism* $m = \langle \mathbf{k}, \mathbf{x} \rangle$ *satisfies interim budget balance if*

$$\mathbb{E}^m \left[\sum_{s=t}^{T} \delta^{s-t} \sum_{i=1}^{N} x_i(\theta_s|h^{s-1}) \mid h^{t-1} \right] \geq 0$$

$\forall\ h^{t-1} \in H^{t-1}$, $\forall t$.

Note that

$$\mathbb{E}^m \left[\sum_{s=t}^{T} \delta^{s-t} \sum_{i=1}^{N} x_i(\theta_s|h^{s-1}) \mid h^{t-1} \right] = \sum_{i=1}^{N} \mathbb{E}^m \left[\sum_{s=t}^{T} \delta^{s-t} u_i(\theta_s|h^{s-1}) - U_i(\theta_t|h^{t-1}) \mid h^{t-1} \right] \tag{2.3}$$

Thus, the definition(s) can be written equivalently in terms of $\langle \mathbf{k}, \mathbf{U} \rangle$.

Finally, if the budget was required to balance in expectation only at the start of the contract for history h^0, then we call the constraint *ex ante budget balance.*

2.6 Efficient Mechanisms

Proposition 4 and Corollary 4 above help us unify a host of results from the existing literature on efficient dynamic mechanism design. Most (if not all) these papers invoke the ex post notion of efficiency. Define

$$S(\theta, k) = \sum_{i=1}^{N} u_i(\theta, k)$$

to be the (static) economic surplus generated by the N-player mechanism design problem.

Definition 13. *Let $\mathbf{k}^*$ be the first-best allocation rule, that is, for all $\theta_t \in \Theta$, $h^{t-1} \in H^{t-1}$ and $\forall t$,*

$$k^*(\theta_t|h^{t-1}) = arg \max_{k \in K} S(\theta_t, k)$$

A mechanism $m = \langle \mathbf{k}, \mathbf{U} \rangle$ or $\langle \mathbf{k}, \mathbf{x} \rangle$ satisfies efficiency if $\mathbf{k} = \mathbf{k}^*$.

Athey and Segal [2007b, 2013] provide the dynamic version of AGV mechanism that satisfies perfect Bayesian incentive compatibility and ex post budget balance but violates individual rationality at some history. Bergemann and Valimaki [2010] present the dynamic analog of the pivot (or VCG) mechanism that satisfies ex post incentive compatibility and individual rationality every period but not budget balance.[9]

In a bilateral trade setting, Myerson and Satterthwaite [1983] showed that it is simultaneous interaction of incentives, participation and budget balance that leads to an impossibility of efficiency result. It is thus interesting to analyze the interaction of these forces in a dynamic model. Athey and Miller [2007], Skrzypacz and Toikka [2013] and Lamba [2014] provide necessary and sufficient conditions for implementation of

[9]Under some assumptions their mechanism also satisfies efficient exit– an agent that stops being pivotal does get any more transfers.

the efficient allocation under ex post (and perfect Bayesian) incentive compatibility, individual rationality and budget balance in a repeated bilateral trade setting. Athey and Miller [2007] consider only iid shocks. While Skrzypacz and Toikka [2013] only look at ex ante budget balance, Lamba [2014] provides results for interim and ex post budget balance. Addition of a savings technology laxes the dynamic budget balance constraint; Lamba [2014] explores this permissibility.

The common thread in all these papers is that they implement the efficient allocation. Hence, all the mechanisms constructed in these papers, while different in their inter-temporal structure of payments, must lie in the same equivalence class of history dependent constants presented in Proposition 4 (or Corollary 4) that implement the efficient allocation.

We offer a simple translation of the dynamic VCG mechanism of Bergemann and Valimaki [2010] into the dynamic collateral VCG mechanism proposed in Lamba [2014] using the dynamic payoff equivalence result. Suppose our objective is to *implement the efficient allocation under (perfect Bayesian) incentive compatibility, individual rationality and interim budget balance.*

We can start with the VCG mechanism. Define

$$S_{-i}(\theta_{-i}, k) = \sum_{j \neq i} u_j(\theta_{-i}, k)$$

to be surplus generated without agent i, and let $S^*(\theta)$ and $S^*_{-i}(\theta_{-i})$ be the respective maximum values (over the allocation rule k). As is standard, the VCG mechanism is then defined by $\langle \mathbf{k}^*, \mathbf{x}^{\mathbf{vcg}} \rangle$, where

$$x_i^{vcg}(\theta_t | h^{t-1}) = S^*(\theta_t) - S^*_{-i}(\theta_{-i,t}),$$

for all $\theta_t \in \Theta$, $h^{t-1} \in H^{t-1}$ and $t \leq T$. Next, consider the isomorphic mechanism in terms of the expected utility vectors, $\langle \mathbf{k}^*, \mathbf{U^{vcg}} \rangle$, defined by

$$U_i^{vcg}(\theta_t|h^{t-1}) = \mathbb{E}\left[\sum_{s=t}^{T} \delta^{s-t}\left(S_{-i}^*(\theta_t) - S_{-i}^*(\theta_{-i,t})\right) \mid h^{t-1}, \theta_t\right],$$

where $S_{-i}^*(\theta_t) = \sum_{j \neq i} u_i(\theta_t, k^*)$. This mechanism is ex post incentive compatible and individually rational but it may violate budget balance; for example, in the bilateral trade setting with private goods.

Using the dynamic payoff equivalence result, it is easy to construct another mechanism $\langle \mathbf{k}^*, \mathbf{U}^* \rangle$ where

$$U_i^*(\theta_t|h^{t-1}) = U_i^{vcg}(\theta_t|h^{t-1}) - \inf_{\theta_{i,t} \in \Theta_i} U_i^{vcg}(\theta_t|h^{t-1})$$

Note that in view of Proposition 4 our construction entails

$$a_i(\theta_{-i,t}|h^{t-1}) = -\inf_{\theta_{i,t} \in \Theta_i} U_i^{vcg}(\theta_t|h^{t-1}).$$

Hence, the new mechanism is ex post incentive compatible. Also, by construction it is individually rational.

This mechanism, which we call the dynamic collateral VCG mechanism essentially transfers all expected surplus at every history to the mechanism designer while leaving the minimalist possible utility for the agents that satisfies incentives and participation. It is then immediate to see that if this mechanism does not satisfy interim budget balance, no other mechanism will.

For all $t \leq T$, and $h^{t-1} \in H^{t-1}$, let $EBS^*(h^{t-1})$ represent the values calculated in equation (2.3) for the dynamic collateral VCG mechanism. Then, we have the following result.

Proposition 5. *There exists an ex post incentive compatible and ex post individually rational mechanism that implements the efficient allocation under interim budget* *balance if and only if* $EBS^*(h^{t-1}) \geq 0 \; \forall h^{t-1}, \forall t$.

Lamba [2014] proves this result for the special case of two player bilateral bargaining and private good in a discrete types model. That paper also shows that $EBS^*(h^{t-1}) \geq 0 \; \forall h^{t-1}, \forall t$ is also a necessary and sufficient condition for implementation under ex post budget balance when considering perfect Bayesian implementability. We conjecture that a similar result can be shown for this more general framework adjusting for measure theoretic considerations.[10]

2.7 Application: Bilateral Trade

The dynamic payoff equivalence result is particularly useful in analyzing second-best mechanisms in standard dynamic mechanism design problems with budget balance. We consider one of the simplest and widely applied problem from static mechanism design of bargaining with two sided asymmetric information.

A buyer and a seller bargain repeatedly. The seller can produce a non-durable good every period at a hidden cost and the buyer has a hidden valuation. Both the cost and the valuation may change over time. In terms of our model, the set of agents is $\{B, S\}$; set of possible valuations and costs are respectively $\Theta_B = \mathcal{V}$ and $\Theta_S = \mathcal{C}$; and the allocation rule is simple the probability of trade, so $k_t = p_t \in [0, 1]$. Stage payoffs for the buyer and the seller are respectively given by[11]

$$v_t p_t - x_{B,t} \quad \text{and} \quad x_{S,t} - c_t p_t$$

The problem of the second-best is then simply the maximization of the expected aggregate surplus or gains from trade subject to incentive compatibility, individual

[10] It will require a repeated use of the dynamic payoff equivalence result to redistribute transfers across time and types.

[11] To keep the outside option normalized to zero, we have tweaked the model a bit to allow for transfers "from" the buyer and transfers "to" the seller.

rationality and interim budget balance constraints.

$$\max_{\langle \mathbf{p},\mathbf{U}\rangle} \quad GFT = \int_{\mathcal{V}}\int_{\mathcal{C}} \Big[(v_1 - c_1)\, p(v_1, c_1) +$$

$$\delta \int_{\mathcal{V}}\int_{\mathcal{C}} (v_2 - c_2)\, p(v_2, c_2|v_1, c_1) g(c_2|c_1) f(v_2|v_1) dc_2 dv_2 \Big] g(c_1) f(v_1) dc_1 dv_1$$

subject to

$$(IC) \quad U_B(v_1) \geq U_B(v_1'; v_1)\ \forall v_1, v_1' \in \mathcal{V} \quad \text{and} \quad U_S(c_1) \geq U_S(c_1'; c_1)\ \forall c_1, c_1' \in \mathcal{C}$$

$$U_B(v_2|h) \geq U_B(v_2'; v_2|h)\ \forall v_2, v_2' \in \mathcal{V} \quad \text{and}$$

$$U_S(c_2|h) \geq U_S(c_2'; c_2|h)\ \forall c_2, c_2' \in \mathcal{C},\ \forall h \in \mathcal{V} \times \mathcal{C}$$

$$(IR) \quad U_B(v_1) \geq 0\ \forall v_1 \in \mathcal{V} \quad \text{and} \quad U_S(c_1) \geq 0\ \forall c_1 \in \mathcal{C}$$

$$U_B(v_2|h) \geq 0\ \forall v_2 \in \mathcal{V} \quad \text{and} \quad U_S(c_2|h) \geq 0\ \forall c_2 \in \mathcal{C},\ \forall h \in \mathcal{V} \times \mathcal{C}$$

and,

$$(BB) \quad EBS = \int_{\mathcal{V}}\int_{\mathcal{C}} \Big[(v_1 - c_1) p(v_1, c_1) - U_B(v_1, c_1) - U_S(v_1, c_1) +$$

$$\delta \int_{\mathcal{V}}\int_{\mathcal{C}} (v_2 - c_2) p(v_2, c_2|v_1, c_1) g(c_2|c_1) f(v_2|v_1) dc_2 dv_2 \Big] g(c_1) f(v_1) dc_1 dv_1 \geq 0$$

$$EBS(h) = \int_{\mathcal{V}}\int_{\mathcal{C}} \Big[v_2 - c_2) p(v_2, c_2|h) - U_B(v_2, c_2|h) - U_S(v_2, c_2)|h \Big]$$

$$g(c_2|h_c) f(v_2|h_v) dc_2 dv_2 \geq 0,\ \forall h = (h_v, h_c) \in \mathcal{V} \times \mathcal{C}$$

The analysis in this section generalizes the approach of Myerson and Satterthwaite [1983] to the dynamic model. We collapse the necessary local incentive compatibility conditions into an envelope formula which is then plugged into the expected budget surplus constraint. Note, however, that because the problem is dynamic this exercise is repeated for every possible history. We first provide the flavor of the general argument, and then proceed to solve specific examples to formally lay down the key forces in our understanding of efficient implementation.

Let $\mathcal{V} = [\underline{v}, \overline{v}]$ and $\mathcal{C} = [\underline{c}, \overline{c}]$.[12] We have the following characterization of expected budget surplus.

Lemma 4. *For any perfect Bayesian incentive compatible mechanism,*

$$EBS + U_B(\underline{v}) + U_S(\overline{c}) = \int\limits_{\mathcal{V}}\int\limits_{\mathcal{C}} \Bigg[\left\{ \left(v_1 - \frac{1 - F(v_1)}{f(v_1)} \right) - \left(c_1 + \frac{G(c_1)}{g(c_1)} \right) \right\} p(v_1, c_1)$$

$$+\delta \int\limits_{\mathcal{V}}\int\limits_{\mathcal{C}} \left\{ \left(v_2 - \frac{1 - F(v_1)}{f(v_1)} \left[-\frac{\partial F(v_2|v_1)/\partial v_1}{f(v_2|v_1)} \right] \right) - \left(c_2 + \frac{G(c_1)}{g(c_1)} \left[-\frac{\partial G(c_2|c_1)/\partial c_1}{g(c_2|c_1)} \right] \right) \right\}$$

$$\times p(v_2, c_2|v_1, c_1) g(c_2|c_1) f(v_2|v_1) dc_2 dv_2 \Bigg] g(c_1) f(v_1) dc_1 dv_1,$$

and

$$EBS(h) + U_B(\underline{v}|h) + U_S(\overline{c}|h) = \int\limits_{\mathcal{V}}\int\limits_{\mathcal{C}} \left[\left(v_2 - \frac{1 - F(v_2|h_v)}{f(v_2|h_v)} \right) - \left(c_2 + \frac{G(c_2|h_c)}{g(c_2|h_c)} \right) \right]$$

$$\times p(v_2, c_2|h) g(c_2|h_c) f(v_2|h_v) dc_2 dv_2$$

It is important to note that the above result can be stated and proven verbatim for ex post incentive compatibility. The key to the indifference is that the expected budget surplus takes expectation over all current and future types of both agents.

Lemma 4 pins down the value of the expected budget surplus in terms of the primitives and the allocation rule up to an additive constant. Using the dynamic payoff equivalence result these additive constants determine all possible values of the

[12] For simplicity, in this section, in a slight abuse of notation, we will denote $U_B(\underline{v}) = \inf_{v \in \mathcal{V}} U_B(v)$ and $U_B(\overline{c}) = \inf_{c \in \mathcal{C}} U_S(c)$. This is true if Markov processes satisfy first-order stochastic dominance, but not in general. It is, however, not essential for our results.

expected budget surplus for an incentive compatible mechanism implementing a fixed allocation.

Let EBS^{**} and $EBS^{**}(h)$ refer specifically to the incentive compatible and individually rational mechanisms where

$$U_B(\underline{v}) = U_S(\overline{c}) = U_B(\underline{v}|h) = U_S(\overline{c}|h) = 0$$

for all $v \in \mathcal{V}$, $c \in \mathcal{C}$, and $h \in \mathcal{V} \times \mathcal{C}$.

Corollary 5. *A perfect Bayesian (or ex post) incentive compatible and individually rational mechanism can be implemented under interim budget balance if and only if* $EBS^{**} \geq 0$ *and* $EBS^{**}(h) \geq 0$ *for all* $h \in \mathcal{V} \times \mathcal{C}$.

Proof. Sufficiency is obvious. For necessity, note that EBS^{**} and $EBS^{**}(h)$ are the highest expected budget surpluses that can be generated for their respective histories while satisfying IC and IR. If they are not non-negative no other IC and IR mechanism can ensure them to be. □

Removing transfers, the second best mechanism can then be explicitly formulated by the following result.

Corollary 6. *A perfect Bayesian incentive compatible and individually rational allocation maximizes expected gains from trade under interim budget balance if only if it solves*

$$\max_{\mathbf{p}} \; GFT$$

subject to

$$EBS^{**} \geq 0, \quad EBS^{**}(h) \geq 0 \;\forall h, \; \textit{and}$$

$$\mathbf{p} \textit{ is PBIC}$$

Implementability of $\mathbf{p}$ is essentially a requirement that global incentive constraints be satisfied. In the static model it is replaced by the familiar monotonicity condition on the allocation. In the dynamic model, Pavan, Segal and Toikka [2014] show that it can instead be replaced by what they call the *integral monotonicity* condition. In the dynamic model though the implications of this condition are much less understood and global incentive constraints can routinely bind; see Battaglini and Lamba [2014].[13]

A substantial appeal of the static model with linear preferences and continuous types is that the solution is always bang-bang, the probability of trade is always 0 or 1. Since the objective and all constraints are linear in the allocation, the same insight goes through in the dynamic model too.

Also, Lamba [2014] shows that if there exists a mechanism that implements an allocation under interim budget balance, then there must exist a mechanism that implements it under ex post budget balance.[14] The argument for a finite period model simply follows an inductive inter-temporal redistribution of transfers.

Perhaps one of the most studied simple expositions of the competing economic forces of information, participation and budget balance in mechanism design is the uniform bilateral trading problem- static version of our model with a uniform prior.[15] In order to clearly bring out the prominent economic forces in a simple fashion, we work in a natural extension of this model to the dynamic environment, and present the optimal mechanisms under the extreme ends of the information space. In the next three subsections the types of the buyer and the seller are assumed to be uniform on [0,1] in both periods.[16]

[13]In most of the literature fairly stringent sufficient monotonic conditions (requiring monotonicity of the allocation across histories) are invoked ex post to check whether the allocation in incentive compatible.

[14]This is valid only for perfect Bayesian incentive compatibility and not ex post.

[15]See Myerson [1985] and Gibbons [1992].

[16]We use the iid model to motivate ideas because it offers a simple and complete characterization. We do not have to worry about global incentive constraints. Though a characterization with persistence for a simple model is an important benchmark to work towards.

2.7.1 Static Benchmark

Following Myerson and Satterthwaite [1983], let us first consider the case where $\delta = 0$. Then, writing down the problem in corollary 2 as a Lagrangian[17], it is easy to show that trade happens, that is, $p(v_1, c_1) = 1$, if and only if $v_1 > c_1 + M$, where M solves the binding EBS constraint

$$\int_0^1 \int_0^1 \mathbf{1}_{\{v_1 > c_1 + M\}} \left[2v_1 - 1 - 2c_1\right] dc_1 dv_1 = 0$$

$$\text{i.e.,} \quad \frac{1}{6}(4M - 1)(1 - M)^2 = 0$$

Thus, trade happens if and only if $v_1 > c_1 + \frac{1}{4}$. Since efficiency demands trade for every $v_1 \geq c_1$; $M = \frac{1}{4}$ precisely characterizes the no trade region and the degree of inefficiency. Since the allocation is monotonic, it is implementable. Figure 2.1a captures the no-trade region pictorially. The solid diagonal represents the locus $v_1 = c_1$. Efficient allocation requests trade above the solid diagonal. The area above the dotted line represents the actual trade region.

2.7.2 IID case

Now, suppose the types of both agents are distributed uniformly on [0,1] in both periods.[18] We first consider implementation under ex ante budget balance. In the second period distortions in the EBS constraint are 0, and trade is always efficient, generating a maximum possible surplus of $\frac{1}{6}$ irrespective of the history in period 1.

[17] $\int_0^1 \int_0^1 \left[(v_1 - c_1) + \lambda \left(2v_1 - 1 - 2c_1\right)\right] p(v_1, c_1) dc_1 dv_1$, where λ is the multiplier on EBS. Details are in the appendix.

[18] Formal proofs are presented in the appendix.

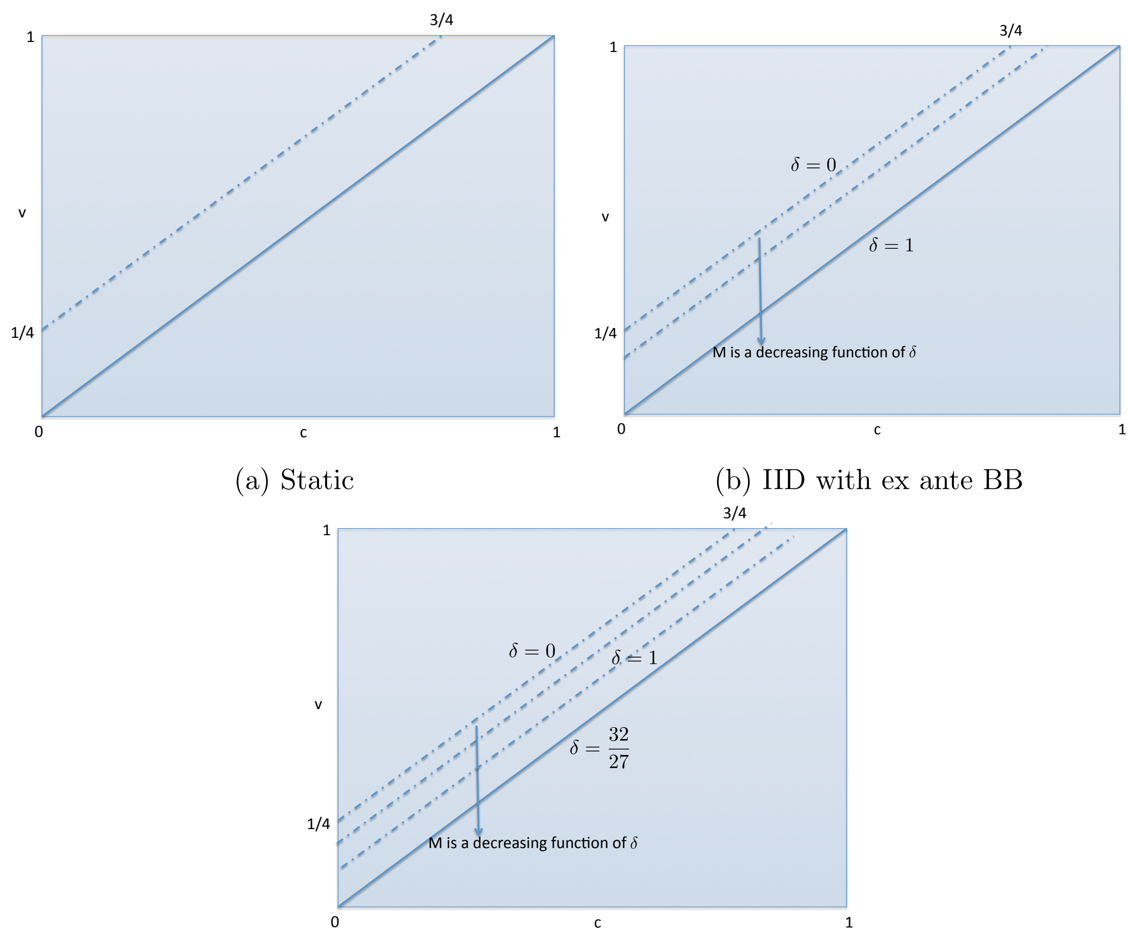

(a) Static

(b) IID with ex ante BB

(c) IID with interim BB

Figure 2.1: No-trade regions

Thus, the no trade region in period 1 solves

$$\frac{1}{6}(4M-1)(1-M)^2+\delta\frac{1}{6}=0 \tag{2.4}$$

It is clear that M is a decreasing function of δ. In fact, at $\delta=0$, $M=1/4$, replicating the static model, and $M=0$ for $\delta=1$, implying the implementability of the efficient allocation.

In the T period version of this problem, trade is always efficient starting period 2. Using a recursive mechanism, Athey and Miller [2007] show that in fact for any distribution, when $T=\infty$, trade in first period will be efficient for any $\delta \geq \frac{1}{2}$.

Next, consider implementation under interim (or ex post budget balance) budget balance. Now, it is easy to show that $EBS^{**}(h)$ will always bind, and the trade in the second period replicates the static model. Thus, $p(v_2, c_2|h)=1$ if and only if $v_2 > c_2 + \frac{1}{4}$ for all h. In the first period, the no trade region M solves

$$\frac{1}{6}(4M-1)(1-M)^2+\delta\frac{9}{64}=0 \tag{2.5}$$

Again, M is a decreasing function of δ. However, this time even $\delta=1$ cannot guarantee efficient trade in period 1. Interim budget balance forces the agents to internalize incentives for period 2 while deciding on the optimal mechanism. The contract is no longer efficient in period 2, and hence a smaller *collateral* is available for trade in period 1.

Nevertheless, it is interesting to consider δ as a proxy for the surplus available in the future in a general T period model. Following the said motivation, it is easy to see that trade will be efficient, and the expected budget surplus constraint will not bind in period 1 for $\delta \geq \frac{32}{27}$.

2.7.3 Perfect Persistence

When types are perfectly persistent, there are no second period expected budget surplus constraints. So, all notions of budget balance are equivalent. The problem reduces to the maximization of gains from trade under ex ante BB and constant types, thus giving us a repetition of the static optimum. In both periods trade happens if and only if $v > c + \frac{1}{4}$.

2.7.4 Discussion

We know that efficient trade is impossible in the static model. What is the driving force in the dynamic model that presents us with the possibility of efficiency? A casual glance at equations (2.4) and (2.5) gives us an indication. A greater expected economic surplus creates a collateral that can used to sustain efficiency. This surplus generated is a function of discounting, the levels of asymmetric information through persistence and the limits on insurance imposed by varying notions of budget balance.[19]

Under iid types and ex ante budget balance, maximum possible surplus is generated in period 2, which can distributed across types and time, significantly reducing the no-trade region in period 1. However, under a stricter notion of budget balance, there are limits on the depth of the credit line facilitating trade. It reduces the total future surplus, thereby mitigating the advantage dynamics present even for the case with minimal informational constraints.

At the other extreme, when the informational constraints are the most severe, in the form of constant types, it blunts all possible advantages that dynamics present by making all those histories that would relax information constraints to generate future surplus be zero in probability.

[19]See Lamba [2014] for a full discussion of efficiency in repeated bargaining.

2.8 Conclusion

This paper introduces a dynamic and history dependent version of the payoff equivalence result. We show that given two mechanisms that implement the same allocation, expected utility of an agent after any history in one must differ from the other through a history dependent constant. This result is then used to unify existing results in efficient dynamic mechanism design. For a general, N- player mechanism design problem, we also precisely characterize the conditions under which the efficient allocation can be implemented under participation constraints and budget balance. This conditions puts joint restrictions on the stochastic evolution of types and the level of discounting. We also show how the dynamic payoff equivalence result can be exploited to calculate the second-best mechanisms through a simple of repeated bilateral trading. Finally, we elucidate the nuances of degree of transparency in the mechanism and its impact on the dynamic payoff equivalence result.

Going forward it will be useful to study specific dynamic mechanism design problems, especially those related to multiple goods auctions and public goods provision where the dynamic equivalence formula will prove very useful in characterizing the revenue and social welfare maximizing mechanisms.

2.9 Appendix

2.9.1 Proof of Proposition 4

Sufficiency. Suppose $\langle \mathbf{k}, \mathbf{U} \rangle$ is ex post incentive compatible. Fix $\theta_{-i,t}$ and h^{t-1}. Then, $U_i(\theta_{i,t}, \theta_{-i,t}|h^{t-1})$ appears in two kinds of incentive compatibility constraints. First,

$$U_i(\theta_{i,t}, \theta_{-i,t}|h^{t-1}) \geq U_i(\theta'_{i,t}, \theta_{-i,t}|h^{t-1}) + \left(u_i(\theta'_{i,t}; \theta_t|h^{t-1}) - u_i(\theta'_{i,t}, \theta_{-i,t}|h^{t-1})\right)$$

$$+\delta \int_{\Theta_i} U_i(\theta_{i,t+1}|h^{t-1}, \theta'_{i,t}, \theta_{-i,t}) \left(f_i(\theta_{i,t+1}|\theta_{i,t}) - f_i(\theta_{i,t+1}|\theta'_{i,t}) \right) d\theta_{i,t+1} \tag{2.6}$$

and,

$$U_i(\theta'_{i,t}, \theta_{-i,t}|h^{t-1}) \geq U_i(\theta_{i,t}, \theta_{-i,t}|h^{t-1}) + \left(u_i(\theta_{i,t}; \theta'_{i,t}\theta_{-i,t}|h^{t-1}) - u_i(\theta_{i,t}, \theta_{-i,t}|h^{t-1}) \right)$$

$$+\delta \int_{\Theta_i} U_i(\theta_{i,t+1}|h^{t-1}, \theta_{i,t}, \theta_{-i,t}) \left(f_i(\theta_{i,t+1}|\theta'_{i,t}) - f_i(\theta_{i,t+1}|\theta_{i,t}) \right) d\theta_{i,t+1} \tag{2.7}$$

Clearly addition of any constant $a_i(\theta_{-i,t}|h^{t-1})$ to $U_i(\theta_{i,t}, \theta_{-i,t}|h^{t-1})$ for all $\theta_{i,t} \in \Theta_i$ does not affect any of these constraints.

Next, fix $(\theta_{i,t-1}, \theta_{-i,t-1}) = h^{t-1}_{t-1}$. Second, we need to consider the constraints,

$$U_i(\theta'_{i,t-1}, \theta_{-i,t-1}|h^{t-2}) \geq U_i(\theta_{i,t-1}, \theta_{-i,t-1}|h^{t-2}) +$$

$$\left(u_i(\theta_{i,t-1}; \theta'_{i,t-1}\theta_{-i,t-1}|h^{t-2}) - u_i(\theta_{i,t-1}, \theta_{-i,t-1}|h^{t-2}) \right)$$

$$+\delta \int_{\Theta_i} U_i(\theta_{i,t}|h^{t-1}) \left(f_i(\theta_{i,t}|\theta'_{i,t-1}) - f_i(\theta_{i,t}|\theta_{i,t-1}) \right) d\theta_{i,t}$$

Thus, addition of the constants $a_i(\theta_{-i,t}|h^{t-1})$ to $U_i(\theta_{i,t}, \theta_{-i,t}|h^{t-1})$ for all $\theta_{i,t} \in \Theta_i$ leads to addition of $a_i(h^{t-1}) = \mathbb{E}_{\theta_{-i,t}}\left[a_i(\theta_{-i,t}|h^{t-1})|h^{t-1}\right]$ to $U_i(\theta_{i,t}|h^{t-1})$ for all $\theta_{i,t} \in \Theta_i$ which drops out of the constraint.

Therefore, linear additions of constants as defined in the proposition preserves incentives.

Necessity. Fix agent i and history $h^{t-1} \in H^{t-1}$. From (2.6) and (2.7), incentive compatibility implies

$$\left(u_i(\theta_t|h^{t-1}) - (u_i(\theta_{i,t}; \theta'_{i,t}, \theta_{-i,t}|h^{t-1}) \right) +$$

$$\delta \int_{\Theta_i} U_i(\theta_{i,t+1}|h^{t-1}, \theta_t) \left(dF_i(\theta_{i,t+1}|\theta_{i,t}) - dF_i(\theta_{i,t+1}|\theta'_{i,t})\right)$$

$$\geq U_i(\theta_t|h^{t-1}) - U_i(\theta'_{i,t}, \theta_{-i,t}|h^{t-1}) \geq$$

$$\left(u_i(\theta'_{i,t}; \theta_t|h^{t-1}) - u_i(\theta'_{i,t}, \theta_{-i,t}|h^{t-1})\right) +$$

$$\delta \int_{\Theta_i} U_i(\theta_{i,t+1}|h^{t-1}, \theta'_{i,t}, \theta_{-i,t}) \left(dF_i(\theta_{i,t+1}|\theta_{i,t}) - dF_i(\theta_{i,t+1}|\theta'_{i,t})\right)$$

Using integration by parts, this can be written as

$$\left(u_i(\theta_t|h^{t-1}) - (u_i(\theta_{i,t}; \theta'_{i,t}, \theta_{-i,t}|h^{t-1})\right) +$$

$$\delta \int_{\Theta_i} \frac{\partial U_i(\theta_{i,t+1}|h^{t-1}, \theta_t)}{\partial \theta_{i,t+1}} \cdot \left[F(\theta_{i,t+1}|\theta_{i,t}) - F(\theta_{i,t+1}|\theta'_{i,t})\right] d\theta_{i,t+1}$$

$$\geq U_B(v_t, c_t|h^{t-1}) - U_B(v'_t, c_t|h^{t-1}) \geq$$

$$(\left(u_i(\theta'_{i,t}; \theta_t|h^{t-1}) - u_i(\theta'_{i,t}, \theta_{-i,t}|h^{t-1})\right) +$$

$$\delta \int_{\Theta_i} \frac{\partial U_i(\theta_{i,t+1}|h^{t-1}, \theta'_{i,t}, \theta_{-i,t})}{\partial \theta_{i,t+1}} \cdot \left[F(\theta_{i,t+1}|\theta_{i,t}) - F(\theta_{i,t+1}|\theta'_{i,t})\right] d\theta_{i,t+1}$$

Since the utility functions and stochastic processes satisfy all standard regularity conditions[20], the usual envelope argument gives us

$$\frac{\partial U_i(\theta_{i,t}, \theta_{-i,t}|h^{t-1})}{\partial \theta_{i,t}} = \frac{\partial u(\theta_{i,t}, \theta_{-i,t}|h^{t-1})}{\partial \theta_i} + \delta \int_{\Theta_i} \frac{\partial U_i(\theta_{i,t+1}|h^{t-1}, \theta_t)}{\partial \theta_{i,t+1}} \cdot \frac{\partial F(\theta_{i,t+1}|\theta_{i,t})}{\partial \theta_{i,t}} d\theta_{i,t+1}$$

A slight modification of the Theorem 1 in Pavan, Segal and Toikka [2014] tells us that this can be done recursively. The modification being an extra variable in conditioning

[20] See sections 2.1 and 3.1 of Pavan, Segal and Toikka [2014].

on expectations, viz $\theta_{i,t}$, since we are using ex post incentive compatibility. Thus[21]

$$U_i(\theta'_{i,t}, \theta_{-i,t}|h^{t-1}) - U_i(\theta''_{i,t}, \theta_{-i,t}|h^{t-1})$$

$$\int_{\theta''_{i,t}}^{\theta'_{i,t}} \left[\frac{\partial u(\theta_{i,t}, \theta_{-i,t}|h^{t-1})}{\partial \theta_i} + \delta \int_{\Theta_i} \frac{\partial U_i(\theta_{i,t+1}|h^{t-1}, \theta_t)}{\partial \theta_{i,t+1}} \cdot \frac{\partial F(\theta_{i,t+1}|\theta_{i,t})}{\partial \theta_{i,t}} d\theta_{i,t+1} \right] d\theta_{i,t} \quad (2.8)$$

The result follows.

2.9.2 Proof of Proposition 5

Sufficiency is obvious. If the Collateral Dynamic VCG mechanism satisfies all the necessary properties, then it is one such desired mechanism.

Next, we establish necessity. Most of the substance follows from Proposition 4 , that is, payoff equivalence. Suppose there exists a set of histories $\mathcal{H}$ of non-zero measure, such that $EBS^*(h) < 0$ for all $h \in \mathcal{H}$.

Consider any other mechanism $\langle \mathbf{p}, \mathbf{U} \rangle$ that is ex post incentive compatible, individually rational and and implements the efficient allocation under interim budget balance. Then, by construction,

$$U_i(\theta_t|h^{t-1}) \geq U_i^*(\theta_t|h^{t-1}) \quad \forall i \; \forall h^{t-1}$$

Recollect from equation (2.3) that expected budget surplus can be written as

$$EBS(h^{t-1}) = \sum_{i=1}^{N} \mathbb{E}^m \left[\sum_{s=t}^{T} \delta^{s-t} u_i(\theta_s|h^{s-1}) - U_i(\theta_t|h^{t-1}) \mid h^{t-1} \right]$$

Thus, if for any history $h^{t-1} \in \mathcal{H}$, $EBS^*(h^{t-1}) < 0$, we must have $EBS(h^{t-1}) < 0$ in

[21]From any standard integrability theorem, see for example Theorem 5.13 in Royden [1968].

$\langle \mathbf{p}, \mathbf{U} \rangle$.

2.9.3 Blind Mechanism

The blind mechanism consists of each agent announcing her/his type to the mechanism every period; however, the agents cannot observe each other announcements. Thus, the information set of an agent in period t consists of her/his own announcements, actual types and the allocation rules up to period $t-1$. An important restriction on the set of mechanisms then is measurability of the transfers to the agents every period with respect to their information sets. Note that the histories in the blind mechanism only consist of own announcements, own types and the allocation rule.

Keeping these in mind, we can restate the the dynamic payoff equivalence result for the blind mechanism as follows.

Proposition 6. *Payoff equivalence holds after every history for the blind mechanism. That is, if* $\langle \mathbf{k}, \mathbf{U} \rangle$ *and* $\left\langle \mathbf{k}, \tilde{\mathbf{U}} \right\rangle$ *are two perfect Bayesian incentive compatible mechanisms that generate utility vectors* $\left(\left(U_i(\theta_{i,t} | h^{blind,t-1}) \right)_{i=1}^{N} \right)_{t=1}^{T}$ *and* $\left(\left(\tilde{U}_i(\theta_{i,t} | h^{blind,t-1}) \right)_{i=1}^{N} \right)_{t=1}^{T}$ *respectively, then, there exists a family of constants* $\left(\left(a_i(h^{blind,t-1}) \right)_{i=1}^{N} \right)_{t=1}^{T}$ *such that*

$$U_i(\theta_{i,t} | h^{blind,t-1}) = \tilde{U}_i(\theta_{i,t} | h^{blind,t-1}) + a_i(h^{blind,t-1})$$

Conversely, if $\langle \mathbf{k}, \mathbf{U} \rangle$ *is ex post incentive compatible, and* $\mathbf{U}$ *and* $\tilde{\mathbf{U}}$ *satisfy the above two equations for a finite family of constants* $\left(\left(a_i(h^{blind,t-1}) \right)_{i=1}^{N} \right)_{t=1}^{T}$*, then* $\left\langle \mathbf{k}, \tilde{\mathbf{U}} \right\rangle$ *is also ex post incentive compatible.*

To operationalize the blind mechanism and note its basic difference from the public one, we explore the collateral dynamic VCG mechanism. In equilibrium, and on truthful histories, the expected utility vectors are now simply functions of past individual reports.

First, we start with the dynamic VCG mechanism as in section 6, and construct the collateral dynamic VCG mechanism, $\langle \mathbf{k}^*, \mathbf{U}^* \rangle$. This of course cannot be a blind mechanism because it is not measurable with respect to the agents' information sets. Thus, as a final step we define the following utility vectors. Let $\theta_i^t = \{\theta_{i,1}, ..., \theta_{i,t}\}$.

$$U_i^{blind}(\theta_{i,t}|\theta_i^{t-1}) = \mathbb{E}_{\theta_{-i}^t}\left[U^*(\theta_t|h^{t-1}) \mid \theta_i^t\right]$$

And, finally using Proposition 6, we define

$$U_i^{blind,*}(\theta_{i,t}|\theta_i^{t-1}) = U_i^{blind}(\theta_{i,t}|\theta_i^{t-1}) - \inf_{\theta_{i,t}\in\Theta_i} U_i^{blind}(\theta_{i,t}|\theta_i^{t-1})$$

where $a(\theta_i^{t-1}) = -\inf_{\theta_{i,t}\in\Theta_i} U_i^{blind}(\theta_{i,t}|\theta_i^{t-1})$. From the perspective of the mechanism designer who observes the public history, define

$$U_i^{blind,*}(\theta_t|h^{t-1}) = U_i^{blind,*}(\theta_{i,t}|\theta_i^{t-1}),$$

$\forall\ \theta_t \in \Theta$, $\forall\ h^{t-1} \in H^{t-1}$. Thus, the blind mechanism amounts to pooling a lot of the incentive constraints across histories.

For all $t \leq T$, and $h^{t-1} \in H^{t-1}$, let $EBS^{blind,*}(h^{t-1})$ represent the values calculated in equation (2.3) for this blind dynamic collateral VCG mechanism. Note that the set of histories for calculating the expected budget surplus is still public because the mechanism designer can observe all the announcements. Then, we can state the following result.

Proposition 7. *There exists a perfect Bayesian incentive compatible and individually rational blind mechanism that implements the efficient allocation under interim budget balance if and only if* $EBS^{blind,*}(h^{t-1}) \geq 0\ \forall h^{t-1}$, $\forall t$.

The associated transfers can be then be constructed using equation,

$$U_i^{blind,*}(\theta_{i,t}|\theta_i^{t-1}) = u_i^{blind,*}(\theta_{i,t}|\theta_i^{t-1}) - x_i^{blind,*}(\theta_{i,t}|h^{t-1}) +$$

$$\delta \mathbb{E}_{\theta_{t+1}} \left[U_i^{blind,*}(\theta_{i,t+1}|\theta_i^{t-1}, \theta_{i,t}) \left| \theta_i^{t-1}, \theta_{i,t} \right. \right]$$

where

$$u_i^{blind,*}(\theta_{i,t}|\theta_i^{t-1}) = \mathbb{E}_{\theta_{-i}^t} \left[u^*(\theta_t|h^{t-1}) \mid \theta_i^t \right]$$

It is important to note that from the perspective of the mechanism designer

$$U_i^{blind,*}(\theta_t|h^{t-1}) \geq U_i^*(\theta_t|h^{t-1}),$$

$\forall\ \theta_t \in \Theta$, $\forall\ h^{t-1} \in H^{t-1}$. Thus, the blind mechanism produces a weakly higher surplus for every history and thus is more likely to implement the efficient allocation.

Corollary 7. *$EBS^{blind,*}(h^{t-1}) \geq EBS^*(h^{t-1})$ $\forall h^{t-1}$, $\forall t$. Thus, if the public mechanism implements the efficient allocation, so does the blind mechanism.*

2.9.4 Proof of Lemma 4

Incentive compatibility in period 2 for any history h gives,

$$(v_2 - v_2')\, p(v_2|h) \geq U_B(v_2|h) - U_B(v_2'|h) \geq (v_2' - v_2)\, p(v_2'|h)$$

The envelope formula for period 2 thus follows,

$$U_B(v_2|h) = U_B(\underline{v}|h) + \int_{\underline{v}}^{v_2} p(\tilde{v}_2|h) d\tilde{v}_2 \tag{2.9}$$

Employing incentive compatibility in period 1 gives

$$(v_1 - v_1')p(v_1) + \delta \int\limits_{\mathcal{V}} U_B(v_2|v_1) \cdot [dF(v_2|v_1) - dF(v_2|v_1')]$$

$$\geq U_B(v_1) - U_B(v_1') \geq$$

$$(v_1 - v_1')p(v_1') + \delta \int\limits_{\mathcal{V}} U_B(v_2|v_1') \cdot [dF(v_2|v_1) - dF(v_2|v_1')]$$

Using (2.9) and integration by parts, gives the dynamic envelope formula,

$$U_B(v_1) = U_B(\underline{v}) + \int\limits_{\underline{v}}^{v_1} \left[p(\tilde{v}_1) + \delta \int\limits_{\mathcal{V}} p(v_2). \left(-\frac{\partial F(v_2|\tilde{v}_1)/\partial \tilde{v}_1}{f(v_2|\tilde{v}_1)} \right) dv_2 \right] dF(\tilde{v}_1) \quad (2.10)$$

Similarly, we get

$$U_S(c_2|h) = U_S(\bar{c}|h) + \int\limits_{c_2}^{\bar{c}} p(\tilde{c}_2|h) d\tilde{c}_2 \quad (2.11)$$

and,

$$U_S(c_1) = U_S(\bar{c}) + \int\limits_{c_1}^{\bar{c}} \left[p(\tilde{c}_1) + \delta \int_{\mathcal{C}} p(c_2). \left(-\frac{\partial G(c_2|\tilde{c}_1)/\partial \tilde{c}_1}{g(c_2|\tilde{c}_1)} \right) dc_2 \right] dG(\tilde{c}_1) \quad (2.12)$$

Now, in period 2, we can write the expected budget surplus as

$$EBS(h) = \int\limits_{\mathcal{V}} \int\limits_{\mathcal{C}} [x_B(v_2, c_2|h) - x_S(v_2, c_2|h)]\, g(c_2|h_c) f(v_2|h_v) dc_2 dv_2$$

$$= \int\limits_{\mathcal{V}} \int\limits_{\mathcal{C}} [(v_2 p(v_2, c_2|h) - U_B(v_2, c_2|h)) - (c_2 p(v_2, c_2|h) + U_S(v_2, c_2|h))]$$

$$g(c_2|h_c) f(v_2|h_v) dc_2 dv_2$$

which using equations (2.9) and (2.11), and integration by parts can be written as

$$EBS(h)+U_B(\underline{v}|h)+U_S(\overline{c}|h)=\int\limits_{\mathcal{V}}\int\limits_{\mathcal{C}}\left[\left(v_2-\frac{1-F(v_2|h_v)}{f(v_2|h_v)}\right)-\left(c_2+\frac{G(c_2|h_c)}{g(c_2|h_c)}\right)\right]$$

$$\times p(v_2,c_2|h)g(c_2|h_c)f(v_2|h_v)dc_2dv_2$$

Similarly, using equations (2.10) and (2.12) in

$$EBS=\int\limits_{\mathcal{V}}\int\limits_{\mathcal{C}}\Big[(v_1p(v_1,c_1)-U_B(v_1,c_1))-(c_1p(v_1,c_1)+U_S(v_1,c_1))$$

$$+\delta\int\limits_{\mathcal{V}}\int\limits_{\mathcal{C}}\left[v_2p(v_2,c_2|v_1,c_1)-c_2p(v_2,c_2|v_1,c_1)\right]g(c_2|h_c)f(v_2|h_v)dc_2dv_2\Big]g(c_1)f(v_1)dc_1dv_1$$

and integrating by parts, we get the desired equality.

2.9.5 Details of the Examples presented in Section 4

Monotonicity. Let $h^t_{v,j}$ denotes buyer's report in period $j\le t$. For $h^t,\widehat{h}^t\in H^t$, $h^t\succeq\widehat{h}^t$ if $h^t_{v,j}\ge\widehat{h}^t_{v,j}$ and $h^t_{c,j}\le\widehat{h}^t_{c,j}$ for all $j\le t$. Then, an allocation $\mathbf{p}$ is *monotonic* if $h^{t-1}\succeq\widehat{h}^{t-1}\Rightarrow p(v_t|h^{t-1})\ge p(v'_t|\widehat{h}^{t-1})$ and $p(c_t|h^{t-1})\le p(c'_t|\widehat{h}^{t-1})$ for all $v_t\ge v'_t$ and $c_t\ge c'_t$.

In the context of our model, Pavan, Segal and Toikka [2014], and Battaglini and Lamba [2014] show that under first-order stochastic dominance, local incentives and monotonicity imply implementability.

In each of the first three examples, we maximize expected gains from trade under the respective expected budget surplus constraints. In all the cases below, the allocation will be monotonic in the sense described above, so will ignore the implementability constraint. In the static model,

$$\max_{\mathbf{p}} \int_0^1 \int_0^1 (v_1 - c_1)\, p(v_1, c_1) dc_1 dv_1$$

subject to

$$\int_0^1 \int_0^1 (2v_1 - 1 - 2c_1)\, p(v_1, c_1) dc_1 dv_1 \geq 0$$

The Lagrangian can then we written as

$$\int_0^1 \int_0^1 [(v_1 - c_1) + \lambda\, (2v_1 - 1 - 2c_1)]\, p(v_1, c_1) dc_1 dv_1$$

$$= (1 + 2\lambda) \int_0^1 \int_0^1 \left(v_1 - c_1 - \frac{\lambda}{1 + 2\lambda}\right) p(v_1, c_1) dc_1 dv_1$$

Since we know that the efficient allocation is not implementable, we must have $\lambda > 0$. Thus, we must have

$$p(v_1, c_1) = \begin{cases} 1 & \text{if } v_1 > c_1 + M \\ 0 & \text{otherwise} \end{cases}$$

where $M = \frac{\lambda}{1+2\lambda}$. Substituting this allocation rule in the binding constraint, gives $M = \frac{1}{4}$. The two period iid problem under ex ante BB can be written as

$$\max_{\mathbf{p}} \int_0^1 \int_0^1 \left[(v_1 - c_1)\, p(v_1, c_1) + \delta \int_0^1 \int_0^1 (v_2 - c_2)\, p(v_2, c_2 | v_1, c_1) dc_2 dv_2 \right] dc_1 dv_1 \tag{2.13}$$

subject to

$$\int_0^1 \int_0^1 \left[(2v_1 - 1 - 2c_1)\, p(v_1, c_1) + \delta \int_0^1 \int_0^1 (v_2 - c_2)\, p(v_2, c_2 | v_1, c_1) dc_2 dv_2 \right] dc_1 dv_1 \geq 0 \tag{2.14}$$

Setting up the Lagrangian, it is easy to see that the optimal mechanism is efficient in period 2 and M in period 1 solves equation (2.4).

Finally, in the two period iid problem under interim BB we maximize (2.13) subject to (2.14) and

$$\int_0^1 \int_0^1 (2v_2 - 1 - 2c_2)\, p(v_2, c_2 | v_1, c_1) dc_2 dv_2 \geq 0,$$

for all v_1, c_1. Setting up the Lagrangian it is easy to see that allocation in period 2 replicates the static model, and period 1 no-trade region solves equation (2.5).

Bibliography

[1] Claude d'Aspremont and Louis-André Grérard-Varet (1979a). "Incentives and Incomplete Information," *Journal of Public Economics*, 11(1): 25-45.

[2] Susan Athey and David Miller (2007). "Efficiency in Repeated Trade with Hidden Valuations," *Theoretical Economics*, 2(3): 299-354.

[3] Susan Athey and Ilya Segal (2007a). "Designing Efficient Mechanisms for Dynamic Bilateral Trading Games," *American Economic Review P&P* , 97(2): 131-136.

[4] Susan Athey and Ilya Segal (2007b). "An Efficient Dynamic Mechanism," *working paper.*

[5] Susan Athey and Ilya Segal (2013). "An Efficient Dynamic Mechanism," *Econometirca, forthcoming.*

[6] Marco Battaglini and Rohit Lamba (2014). "Optimal Dynamic Contracting: the First-Order Approach and Beyond," *working paper.*

[7] Kim-Sau Chung and Jeff Ely (2003). "Ex-post Incentive Compatible Mechanism Design," *working paper.*

[8] Jacques Crémer and Richard P. McLean (1988). "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56(6): 1247-1257.

[9] Robert Gibbons (1992). *Game Theory for Applied Economists,* Princeton University Press.

[10] Theodore Groves (1973). "Incentives in Teams," *Econometrica*, 41(4): 617-631.

[11] Vijay Krishna and Motty Perry (2000). "Efficient Mechanism Design,"*working paper.*

[12] Rohit Lamba (2014). "Efficiency in Repeated Bargaining: A Mechanism Design Approach," *working paper.*

[13] Mailath, G. and A. Postelwaite (1991). "Asymmetric Information Bargaining Problems with Many Agents," *Review of Economic Studies*, 57(3), 351-367.

[14] Andreu Mas-Colell, Michael Whinston, and Jerry Green (1995). *Microeconomic Theory,* Oxford University Press.

[15] Paul Milgrom (2004). *Putting Auction Theory to Work,* Cambridge University Press.

[16] Roger Myerson (1981). "Optimal Auction Design," *Mathematics of Operations Research*, 6(1): 58-73.

[17] Roger Myerson and Mark Satterthwaite (1983). "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29(2): 265-281.

[18] Roger Myerson (1985). "Analysis of Two Bargaining Problems with Incomplete Information," *Game Theoretic Models of Bargaining*, edited by Al Roth, Cambridge University Press.

[19] Alessandro Pavan, Ilya Segal and Juuso Toikka (2014). "Dynamic Mechanism Design: A Myersonian Approach," *Econometirca,* 82(2): 601-653.

[20] Andy Skrzypacz and Juuso Toikka (2013). "Mechanisms for Repeated Bargaining," *working paper*.

[21] William Vickrey (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16(1): 8-37.

[22] Robert Wilson (1992). *Non-linear Pricing,* Oxford University Press.

Chapter 3

Optimal Dynamic Contracting: the First-Order Approach and Beyond

3.1 Introduction

Most contractual relationships have a dynamic nature, involving long-term, non-anonymous interaction between a principal and an agent. Examples of these contractual relationships include income taxation, regulation, managerial compensation or a monopolist repeatedly selling a non-durable good to a buyer. In these environments contracts can be made contingent on past realizations of the agent's type, allowing the principal to use the agent's revealed preferences to screen future types' realizations. This may be particularly useful in limiting asymmetric information and agency problems when the agent's type is persistent over time.

Despite recent advances in contract theory, there is still a limited understanding about how to use this information to design optimal screening contracts. Dynamic contracts are difficult to study because they involve a large number of incentive compatibility constraints. The analysis of optimal dynamic contracts has therefore been limited to economic environments in which a form of the "first-order approach" can be applied: environments in which the optimal contract can be fully characterized using only the necessary conditions implied by local incentive compatibility constraints. While the first-order approach can be generally applied in static environments under mild regularity assumptions, in dynamic models local incentive compatibility constraints have proven to be sufficient only in very special economic environments.[1]

This leaves three sets of open questions. First, what is the general applicability of the first-order approach and what are its implications? Second, in environments in which the first-order approach does not hold, what does the optimal contract look like? Are there phenomena associated with dynamic contracts that we are ignoring by focusing on environments in which solving the contract is easy? Finally, if characterizing the optimal contract is complicated, can we approximate the optimal contracts with simpler contracts which guarantee a minimal loss in profits?

To address these questions, we consider a simple principal-agent model in which a monopolist repeatedly sells a non durable good to a buyer. The "type" of the buyer that parametrizes his utility is private information, and it evolves over time according to a general N-state Markov process. Higher types are assumed to have higher marginal valuations and their associated conditional distribution on future types first-order stochastically dominates the distribution of lower types.

We present four sets of results. We start by exploring the applicability of the first-order approach. We show that if we ignore global constraints, necessary local incentive compatibility constraints allow us to state a "dynamic envelope theorem" with discrete types through which the agent's equilibrium rent can be expressed just as a function of the expected allocation. In the relaxed problem that only includes local incentive constraints, the dynamic envelope theorem allows a simple characterization of the profit maximizing contract. In keeping with the terminology from the static literature, this contract is referred to as the first-order optimal contract, or FO-optimal

[1]We will discuss the literature in greater detail in Section 8.

contract. We also show that the envelope formula and a simple form of monotonicity of the allocation are sufficient for implementability.[2] Monotonicity requires that if $h^t \succeq \widehat{h}^t$, then $q(h^t) \geq q(\widehat{h}^t)$, where $q(h^t)$ (resp., $q(\widehat{h}^t)$) is the quantity allocated following a history h^t (resp., $\widehat{h}^t$).[3] This condition is only sufficient and quite strong, but it is verified for virtually all environments in which the optimal dynamic contract has been characterized in the existing literature.[4] Although various characterizations of the envelope conditions have been presented over time,[5] this paper is the first to provide a general characterization of the formula and its implications for discrete types. This approach has two advantages. First, most applied works using numerical methods to study dynamic contracts rely either on the discrete type assumption or discrete approximations of the continuous type model. The formula with discrete types presented in this paper, thus, allows an exact characterization. Second, focusing on discrete types allows us to avoid the measure theoretic complications of the case with continuous types which may obscure otherwise simple economic intuitions.

Second, we show that the environments for which the dynamic envelope formula is sufficient to characterize the optimal dynamic contract are very special. With more than two types, when types' correlation is sufficiently high, the first-order optimal contract is generically non-monotonic. Because of this, global constraints are generically binding if the time horizon is sufficiently important (that is when types' persistence, number of periods T and the discount factor δ are high enough). Numerical calculations show that, in general, the level of persistence needed for the failure of the first-order approach is in fact quite low.

---

[2]An allocation is implementable if there exist transfers such that the contract is incentive compatible.

[3]A history is a vector of reports $h^t = (h_1^t, ..., h_t^t)$, so $h^t \succeq \widehat{h}^t$ if $h_j^t \geq \widehat{h}_j^t \ \forall j \leq t$.

[4]Necessary and sufficient conditions for the optimality of the FO-contract can easily be stated, see Section 4. But, these tend to be less intuitive.

[5]See, among others, Baron and Besanko [1984], Besanko [1985], Laffont and Tirole [1996], Courty and Li [2000], Battaglini [2005], Eso and Szentes [2007], Pavan, Segal and Toikka [2014].

These findings on the limits of the first-order approach have important implications for applied work. In many applications of dynamic principal-agent models (including the study of optimal taxation), the key variable for which agents have private information is their income. Using a recent large data set Guvenen et al. [2013a] and [2013b] [6] show that individual income in the U.S. is very persistent and the empirical distribution of income changes has extremely high kurtosis. Therefore, in all applications where income is the key variable it is appropriate to assume that types are highly persistent. These are precisely the environments where our results suggest that the use of the first-order approach is particularly problematic: either it does not work, due to the violation of some global incentive constraint; or it works only because a non-generic stochastic process has been assumed.

Our third contribution is to fully characterize the optimal contract in a simple environment with three types and two periods. The characterization shows that the seller typically finds it optimal to offer a continuation utility in the second period that is not monotonic in the revealed first period type. The optimal contract is characterized by separation of types even when separation is not optimal in static contracts. It is also characterized by what we call *dynamic pooling*: strategic state contingent treatment of types in which types may be initially separated, to be then pooled conditioned on particular histories.

In our final contribution, we make a first step in addressing the problem of designing optimal contract in complex environments with large T and N. We identify a particular class of allocations for which the optimal implementable contract, which we term *monotonic contracts*, can be easily characterized. Quantities in monotonic contracts *are forced* to be non-decreasing in types (a restriction, following Roger Myerson's original terminology, we call *ironing*). Restricting to this subset of contracts is not optimal in general. We show, however, that the optimal monotonic

[6]The dataset consists of a random sample of 10% of the population between the ages of 25 and 60 from 1978 to 2011.

contract converges in probability to the optimal contract as types become highly persistent, or discounting converges to one, or both persistence and discounting converge to one, independent of the order of these limits. In these cases, the loss in the monopolist's profit goes to zero. Further, numerical calculations show that the optimal monotonic contract performs very well, ensuring a minimal loss in objective, even with lower levels of persistence.

We proceed as follows. In Section 3.2 we present the model. In Section 3.3 we present the dynamic envelope formula and the first-order optimal contract. In Section 3.4 we characterize the validity of the first-order approach. In Section 3.5 we establish the limits of the first-order approach in the form of an impossibility result for a general class of dynamic models. In Section 3.6 we completely characterize a three type, two period model. In Section 3.7 we introduce and explore monotonic contracts. In Section 3.8 we provide an overview of the literature. Finally, conclusions are presented in Section 3.9. Proofs can be found in the appendices.

3.2 Model

There are two players, a buyer (or consumer) and a seller (or monopolist). The buyer repeatedly buys a non-durable good from the seller. Consumer of type θ_t enjoys a per-period utility $u(\theta_t, q) - p$ for q units of the good bought at a price p. In every period, the seller produces the good with a cost function $c(q)$. The utility and cost functions satisfy the usual conditions. The utility function $u(\theta_t, q)$ is increasing and differentiable in both arguments, with $u(\theta_t, 0) = 0$; it is concave in q; and it satisfies the single crossing condition:

Assumption 1. *$u_{\theta q}(\theta, q) > 0$ for any θ and q.*

The cost function $c(q)$ is increasing, convex and differentiable with $c'(0) = 0$ and $\lim_{q\to\infty} c'(q) = \infty$. For future reference, let $s(\theta, q) = u(\theta, q) - c(q)$ be the instantaneous surplus generated by a contract that supplies quantity q to a buyer of type θ. In what follows, $s_q(\theta, q)$ and $u_q(\theta, q)$ denote the derivatives with respect to q. To illustrate some of the results, we will repeatedly use the classic version of this model proposed by Mussa and Rosen [1978] in which $u(\theta_t, q) = \theta_t q$ and $c(q) = (1/2)q^2$.

The type θ_t evolves over time according to a Markov process. There are $N+1$ possible types, $\Theta = \{\theta_0, \theta_1, ..., \theta_N\}$, with $\theta_i - \theta_{i+1} = \Delta\theta > 0$ for any $i = 0, ..., N-1$. Let $\mathcal{N} = \{0, 1, 2, ..., N\}$ denote the set of all indices of types, noting that the indices uniquely identify the types. The probability that type k is reached next period if the agent's current type is i is given by $f(\theta_k|\theta_i) = \alpha_{ik}$. Let F be the conditional CDF, defined $F(\theta_j|\theta_i) = \sum_{k=0}^{N-j} \alpha_{i(j+k)}$. The distribution of types conditional on being type i is denoted $\alpha_i = (\alpha_{i0}, \alpha_{i1}, ..., \alpha_{iN})$, where we assume that α_i has full support (so $\alpha_{ij} > 0$ for any i, j), and α_i first-order stochastically dominates α_j for any i and any $j > i$. Given that higher indices imply lower values, first-order stochastically dominance can be stated as:

Assumption 2. *$F(\theta_j|\theta_i) \leq F(\theta_j|\theta_k)$ for any j and $i \leq k$.*

In each period the consumer observes the realization of his own type; the seller, in contrast, can only observe past allocations. At date 0 the seller has a prior $\mu = (\mu_0, ..., \mu_N)$ on the agent's type. For convenience in most of what follows we assume the prior has full support, so $\mu_i > 0$ for any i. This assumption is made only to simplify notation and is not necessary for the results.

In static models, standard concavity assumptions on the objectives and distributional assumptions like monotone hazard rate on the prior ensure the validity of the first-order approach, see for example Stole [2001]. We require the former assumption, but we do not need the latter. Define:

$$\Phi(\theta_i, q) = s(\theta_i, q) - \frac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \cdot [u(\theta_{i-1}, q) - u(\theta_i, q)].$$

Assumption 3. Φ *is concave and has a a unique interior maximum over* q *for all* i.

This assumption rules out situations in which even in the static model the optimal solution is the zero supply corner solution.

We assume that time is discrete and the relationship between the buyer and the seller lasts for $T \geq 2$ periods. In period 1 the seller offers a supply contract to the buyer. The buyer can reject the offer or accept it; in the latter case the buyer can walk away from the relationship at any time $t \geq 1$ if the expected continuation utility offered by the contract falls below the reservation value $\underline{U} = 0$. In line with the standard model of price discrimination, the monopolist commits to the contract that is offered. The common discount factor is $\delta \in (0, 1)$.

It is easy to show that in this environment a form of the revelation principle is valid, which allows us to consider, without loss of generality, only contracts that depend on the history of type revelations, i.e., the contract can be written as $\langle \mathbf{p}, \mathbf{q} \rangle = \left(p\left(\theta_t \left| h^{t-1}\right.\right), q\left(\theta_t \left| h^{t-1}\right.\right)\right)_{t=1}^{T}$, where h^{t-1} and θ_t are, respectively, the public history up to period $t-1$ and the type revealed at time t.[7] In general, h^t can be defined recursively as $h^t = \{h^{t-1}, \theta_t\}$, $h^0 = \emptyset$. The set of possible histories at time t is denoted H^t (for simplicity $H = H^T$). Let κ_t be the mapping that projects the first t elements of a vector. The set of full histories that follow h^t till time t is given by $H(h^t) = \{h \in H | \kappa_t(h) = h^t\}$. It is also useful to define the set $\widehat{H}(h^t) = \{h \in H(h^t) | h_\tau < \theta_0 \ \forall \tau > t\}$. This is the set of histories following h^t in which all realizations after t are lower than θ_0, the highest type.

A strategy for a seller consists of offering a direct mechanism $\langle \mathbf{p}, \mathbf{q} \rangle$ as described above. The strategy of a consumer is, at least potentially, contingent on a richer history $h_A^t = \left\{h_A^{t-1}, \theta_t, \widehat{\theta}_{t-1}\right\}$, where θ_t is the actual type every period and $\widehat{\theta}_t$ is the revealed type. Note that $h_A^0 = \theta_1$. For a given contract, a strategy for the consumer

[7]When it does not create confusion, the subscript of θ signifies time period or a specific type depending on the context: so when we write θ_i, we mean $\theta_{i,t}$. The extra notation would be superfluous in most of the paper.

is simply a function that maps a history h_A^t into a revealed type: $h_A^t \mapsto s(h_A^t)$.

3.3 The first-order approach and the dynamic envelope formula

In this section we characterize the seller's problem and discuss the standard approach that has been used in the literature to solve it: the so called first-order approach. The seller's problem consists of choosing a contract $\langle \mathbf{p}, \mathbf{q} \rangle$ that maximizes profits under two sets of constraints: incentive compatibility constraints, which guarantee that an agent of type i does not want to report being a type j after any history h^t, and individual rationality constraints, which guarantee that all types expect to receive at least their reservation utility $\underline{U} = 0$ after any history h^t. Since the choice of prices and quantities corresponds to the choices of utilities and quantities for the buyer, this problem can be conveniently represented as a choice of $\langle \mathbf{U}, \mathbf{q} \rangle = \left(U\left(\theta_t \left| h^{t-1}\right.\right), q\left(\theta_t \left| h^{t-1}\right.\right)\right)_{t=1}^{T}$, where $U\left(\theta_t \left| h^{t-1}\right.\right)$ is the expected utility of a type θ_t after history h^{t-1}.

The generic incentive compatibility constraint $IC_{i,j}(h^{t-1})$ requires $U(\theta_i|h^{t-1}) \geq U(\theta_j;\theta_i|h^{t-1})$, where $U(\theta_j;\theta_i|h^{t-1})$ is the expected utility of a type θ_i reporting to be a type θ_j at time t after history h^{t-1}. This constraint can be easily rewritten in terms of $\langle \mathbf{U}, \mathbf{q} \rangle$ as:

$$\begin{aligned} U(\theta_i|h^{t-1}) \;\; &\geq \;\; U(\theta_j|h^{t-1}) + \delta \sum_{k=0}^{N} \left(\alpha_{ik} - \alpha_{jk}\right) U(\theta_k|h^{t-1}, \theta_j) \\ &\quad + u(\theta_i, q(\theta_j|h^{t-1})) - u(\theta_j, q(\theta_j|h^{t-1})). \end{aligned} \tag{3.1}$$

The individual rationality constraint for type i at history h^{t-1}, $IR_i(h^{t-1})$, is a simple non-negativity constraint:

$$U(\theta_i|h^{t-1}) \geq 0. \tag{3.2}$$

For future reference, we call *local downward constraints* the incentive constraints that are associated with a deviation to a contiguous lower type (i.e. $IC_{i,i+1}(h^{t-1})$), and *local upward constraints* the incentive constraints that are associated with a deviation to a contiguous higher type (i.e. $IC_{i+1,i}(h^{t-1})$). We refer to all the other constraints as *global*. A contract that satisfies all incentive and individual rationality constraints is said to be *implementable*.

The monopolist's problem is to maximize expected surplus net of the buyer's expected equilibrium rents:

$$\max_{\langle \mathbf{U}, \mathbf{q} \rangle} \left\{ \begin{array}{c} E\left[S(\mathbf{q})\right] - \sum_{i=0}^{N} \mu_i U(\theta_i | h^0) \\ \text{s.t. } \mathbf{q} \geq 0 \text{ and } IR_i(h^{t-1}),\, IC_{i,j}(h^{t-1}) \\ \text{for any } i, j, t \text{ and } h^{t-1} \in H^{t-1}. \end{array} \right\} \tag{3.3}$$

This is a standard maximization problem of a concave function under a system of non-linear constraints. As T and N increase the number of variables and constraints becomes prohibitively large making (3.3) analytically intractable.

The typical approach in the literature is to first study a relaxed problem in which only individual rationality constraint of the lowest type and the local downward constraints $IC_{i,i+1}(h^t)$ are considered. The remaining constraints can be verified ex-post after the solution of the relaxed problem has been characterized.

Definition 1. *A contract is first-order optimal if and only if it maximizes profits under the following constraints:* $IR_N(h^{t-1})$ *and* $IC_{i,i+1}(h^{t-1})$, $\forall i \in \mathcal{N} \backslash \{N\}$, $\forall h^{t-1} \in H^{t-1}, \forall t$.

Interest in FO-optimal contracts is based on the fact that in many environments they coincide with the optimal contracts. Under standard assumptions, the FO-optimal contract coincides with the optimal contract in a static environment, both with finite and continuous type spaces (Stole [2001]).[8] This approach has also been used in all papers that have extended the principal-agent model to dynamic environments: the first-order autoregressive environment (Besanko [1985]) and the Markov environment with two types (Battaglini [2005]). Perhaps, more importantly, the applied literature often focuses on FO-optimal contracts even in the absence of an explicit proof that local incentive constraints are sufficient for implementability.[9]

It is easy to show that when we consider the relaxed problem with only local downward constraints, the incentive compatibility constraints can be assumed to hold as equalities.[10] This allows us to eliminate utilities from the optimization problem and drastically simplify the constraint set. Let us define:

$$\Delta F\left(\theta_j\left|\theta_i\right.\right) = F\left(\theta_j\left|\theta_i\right.\right) - F\left(\theta_j\left|\theta_{i-1}\right.\right).$$

It denotes the effect on the conditional distribution of a marginal change in type in the previous period. It is important to note that first-order stochastic dominance implies $\Delta F\left(\theta_j\left|\theta_i\right.\right) \geq 0$, for all i and j. Recalling that $\widehat{H}\left(h^t\right)$ is the set of histories following h^t in which all realizations after t are lower than θ_0, and representing by h_k the kth element of history h, we have the following characterization of the agent's utility only as a function of $\mathbf{q}$:[11]

[8]A sufficient condition for the FO-optimal contract to be optimal in a static environment is that the prior μ satisfies the monotone hazard rate condition and $u_\theta(\theta, q)$ is not increasing in θ- conditions satisfied, for example, by a uniform prior and $u\left(\theta, q\right) = \theta q$. See Stole [2001] for discussion of these results.

[9]See Section 8 for a discussion of this literature.

[10]The details of the statements made in this section are formally proven in the appendix.

[11]To interpret (3.4), note that, given a history $\hat{h} = (\hat{h}_1, ..., \hat{h}_s)$, $\hat{h}_\tau$ is the realization of the type at time $\tau \leq s$. It follows that $q(\hat{h}_\tau|\hat{h}^{\tau-1})$ is the quantity at time τ when the realized history is $\hat{h}^{\tau-1}$.

Lemma 1. *Corresponding to a FO-optimal contract, we have:*

$$\frac{U(\theta_i|h^{t-1}) - U(\theta_{i+1}|h^{t-1})}{\Delta\theta} = \frac{\int_{\theta_{i+1}}^{\theta_i} u_\theta(x, q(\theta_{i+1}|h^{t-1}))dx}{\Delta\theta} + \sum_{\hat{h}\in\widehat{H}(h^{t-1},\theta_{i+1})} \sum_{\tau>t} \delta^{\tau-t} \left[\begin{array}{c} \prod_{k=t+1}^{\tau} \Delta F\left(\hat{h}_k \middle| \hat{h}_{k-1}\right) \\ \cdot \frac{\int_{\hat{h}_\tau}^{\hat{h}_\tau+\Delta\theta} u_\theta(x,q(\hat{h}_\tau|\hat{h}^{\tau-1}))dx}{\Delta\theta} \end{array} \right] \tag{3.4}$$

for any $i \in \mathcal{N}\backslash\{N\}$, $h^t \in H^{t-1}$ and $t = 1, ..., T$.

Lemma 1 presents a straightforward dynamic extension of the envelope formula introduced by Myerson [1981]. This can be seen by taking δ to zero, in which case the second term on the right hand side vanishes and (3.4) coincides with the classic static formula. The formula in (3.4) allows us to express the marginal rent of a type exclusively as a function of the allocation $\mathbf{q}$.[12] Although the formula is a complicated function of the conditional probabilities and the allocation, in specific environments it is quite tractable.

Example 1. When types are i.i.d. we have $f(\theta_i|\theta_j) = f(\theta_i|\theta_k)$ for all i, j, k, so for all histories $\Delta F\left(\hat{h}_k \middle| \hat{h}_{k-1}\right) = 0$. It follows that $U(\theta_i|h^{t-1}) - U(\theta_{i+1}|h^{t-1}) = \int_{\theta_{i+1}}^{\theta_i} u_\theta(x, q(\theta_{i+1}|h^{t-1}))dx$. If we assume $u(\theta, q) = \theta q$, then $u_\theta(x, q(\theta_{i+1}|h^{t-1})) = q(\theta_{i+1}|h^{t-1})$. It follows that

$$\left[U(\theta_i|h^{t-1}) - U(\theta_{i+1}|h^{t-1})\right] / \Delta\theta = q(\theta_{i+1}|h^{t-1}).$$

[12] A continuous type version of the formula is presented in Baron and Besanko [1984] for the case in which $T = 3$ and in Besanko [1985] for an infinite horizon model with first-order autoregressive types in which shocks have independent realizations. Battaglini [2005] states the formula for a Markov process with two states: (3.4) is a direct, but more involved extension of this result for the case with $|\Theta| \geq 2$. Pavan, Segal and Toikka [2014] present a general version of the formula for a continuous type space and other stochastic processes.

In particular this holds for the null history, h^0. Thus, the expected rent at $t = 1$ depends only on quantities in the first period and is same as in the static model. The agent has no private information about future realizations beyond period 1 when the contract is signed. So he or she is unable to extract any rents for $t \geq 2$.

Example 2. Assume $u(\theta, q) = \theta q$ and, as in Baron and Besanko [1984], that types are constant, i.e. $f(\theta_i | \theta_i) = 1$ for all $i = 0, ..., N$. In this case, after (h^{t-1}, θ_{i+1}), only history $\hat{h} = \{h^{t-1}, \theta_{i+1}..., \theta_{i+1}\}$ (in which the type remains equal to θ_{i+1}) has positive probability and $\Delta F(\theta_i | \theta_i) = 1$ for all i. Applying (3.4), it follows that:

$$\left[U(\theta_i|h^{t-1}) - U(\theta_{i+1}|h^{t-1})\right]/\Delta\theta = \sum_{\tau \geq t} \delta^{\tau - t} \cdot q(\hat{h}_\tau|\hat{h}^{\tau-1})$$

for all $i \in \mathcal{N} \backslash \{N\}$, where $\hat{h} \in H(h^{t-1}, \theta_{i+1})$ is the history that has all realizations following period t equal to θ_{i+1}. The expected rents are thus a discounted sum of quantities along the constant histories.

Example 3. Assume $u(\theta, q) = \theta q$ and two types, $\theta_0 = \theta_H$ and $\theta_1 = \theta_L$ that are imperfectly correlated. In this case all histories except the "lowest history" (in which all the types' realizations are always θ_L) disappear from (3.4). Given this, we obtain:

$$\left[U(\theta_H|h^{t-1}) - U(\theta_L|h^{t-1})\right]/\Delta\theta = \sum_{\tau \geq t} \delta^{\tau - t} \cdot [F(\theta_L|\theta_L) - F(\theta_L|\theta_H)]^{\tau - t} \cdot q(\widehat{h}_\tau|\widehat{h}^{\tau-1})),$$

where $\widehat{h} = \{h^{t-1}, \theta_L..., \theta_L\}$ is the history following h^{t-1} in which all realization after $t-1$ are θ_L. In this case the rent of the agent at $t = 1$ depends only on the quantities in the lowest history, in which the realizations are always θ_L. This is the envelope formula derived in Battaglini [2005].

Example 4. Another example that will prove useful in the remainder of the paper is when $T = 2$. Assuming the usual utility $u(\theta, q) = \theta q$, the rents at $t = 2$ are given by $U(\theta_i|h^1) - U(\theta_{i+1}|h^1)/\Delta\theta = q(\theta_{i+1}|h^1)$ and those at $t = 1$ by $[U(\theta_i) - U(\theta_{i+1})]/\Delta\theta =$

$q(\theta_{i+1}) + \sum_{k=1}^{N} \delta \Delta F\left(\theta_k|\theta_{i+1}\right) \cdot q(\theta_k|\theta_{i+1})$.

Returning to the general model, we can express the utility vector solely as a function of $\mathbf{q}$ using Lemma 1. Define:

$$U^*(\theta_i|h^{t-1};\mathbf{q}) = \sum_{n=1}^{N-i} \left[\begin{array}{c} \int_{\theta_{i+n}}^{\theta_{i+n-1}} u_\theta(x, q(\theta_{i+n}|h^{t-1}))dx \\ + \sum_{\hat{\mathbf{h}} \in \widehat{H}(\mathbf{h}^{t-1},\theta_{i+n})} \sum_{\tau>t} \delta^{\tau-t} \left[\begin{array}{c} \prod_{k=t+1}^{\tau} \Delta F\left(\hat{h}_k \middle| \hat{h}_{k-1}\right) \\ \cdot \int_{\hat{h}_\tau}^{\hat{h}_\tau+\Delta\theta} u_\theta(x, q(\hat{h}_\tau|\hat{h}^{\tau-1}))dx \end{array} \right] \end{array} \right] \tag{3.5}$$

for any $i < N$, and $U^*(\theta_N|h^{t-1};\mathbf{q}) = 0$. Corollary 1, thus, immediately follows from (3.4):

Corollary 1. *Corresponding to a FO-optimal contract, we have* $U(\theta_i|h^{t-1}) = U^*(\theta_i|h^{t-1};\mathbf{q})$ *for any* $i \in \mathcal{N}$, $h^{t-1} \in H^{t-1}, \forall t$.

The FO-optimal contract can now be characterized as the solution of the following program:

$$\max_{\mathbf{q}\geq 0} \left\{ E\left[S(\mathbf{q})\right] - \sum_{i=0}^{N} \mu_i U^*(\theta_i|h^0;\mathbf{q}) \right\} \tag{3.6}$$

This problem can be solved to obtain the closed form solution. Let $D(h^t)$ be equal to 1 at $t = 1$, and for $t > 1$, define:

$$D(h^t) = \begin{cases} 0 & \text{if } h^t_\tau = \theta_0 \text{ for any } \tau \leq t \\ \prod_{\tau=1}^{t-1} \left(\frac{\Delta F(h^t_{\tau+1}|h^t_\tau)}{f(h^t_{\tau+1}|h^t_\tau)} \right) & \text{else} \end{cases} \tag{3.7}$$

These are the dynamic distortions associated with the FO-optimal contract. Recall that for any θ_i, $s(\theta_i, q)$ is the per period surplus (i.e. $u(\theta_i, q) - c(q)$), and $s_q(\theta_i, q)$ its derivative with respect to q. From the first-order necessary conditions of (3.6) we can easily characterize the FO-optimal contract as follows.[13]

[13]Note that, in the following expression, $D(h^{t-1}, \theta_i)$ corresponds to $D(h^t)$ for $h^t = \{h^{t-1}, \theta_i\}$. Also, θ_{-1} is any dummy type.

Proposition 1. *Corresponding to a FO-optimal contract we have:*

$$s_q(\theta_i, q^*\left(\theta_i|h^{t-1}\right)) \leq \frac{1-\sum_{k=j}^{N}\mu_k}{\mu_j} \cdot D(h^{t-1}, \theta_i) \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta q}(x, q(\theta_i|h^{t-1}))dx \tag{3.8}$$

for any $i \in \mathcal{N}$, $h^{t-1} \in H^{t-1}$ and t, where $\theta_j = h_1^t$, and the above is satisfied with equality if $q^*\left(\theta_i|h^{t-1}\right) > 0$.

It is customary in the literature to assume that the objective function in (3.6) is concave (see Stole [2001] for example): in this case (3.8) is necessary and sufficient and so it uniquely defines the FO-optimal contract. Although this assumption is not required for the following results, it is always verified if we assume preferences a' la Mussa and Rosen [1978], when $u(\theta, q) = \theta q$ and $c(q) = (1/2)q^2$. In this case, at an interior solution, we have:

$$q^*\left(\theta_i|h^{t-1}\right) = \theta_i - \frac{1-\sum_{k=j}^{N}\mu_k}{\mu_j} D(h^{t-1}, \theta_i)\Delta\theta \tag{3.9}$$

where $\theta_j = h_1^t$. Under Assumption 3, moreover, the objective function (3.6) is concave if types are sufficiently persistent.

We can now apply (3.8) to the examples discussed above.

Example 1 (cont.). From (3.7) and (3.8) we can see that when types are i.i.d., it is optimal to offer the optimal static contract in the first period and the efficient contract in all following periods since the quantities offered after $t = 1$ do not affect rents. For the standard model, we have $q^*\left(\theta_i\right) = \theta_i - \frac{1-\sum_{k=i}^{N}\mu_k}{\mu_i}\Delta\theta$ in the first period and $q^*\left(\theta_i|h^{t-1}\right) = \theta_i$ in the following periods.

Example 2 (cont.). From (3.9), it follows that when types are constant it is optimal to offer the same quantities $q^*\left(\theta_i\right) = \theta_i - \frac{1-\sum_{k=i}^{N}\mu_k}{\mu_i}\Delta\theta$ in all periods, irrespective of the history of types' realizations. To see this, note that on histories in which types remain constant we have $D(h^{t-1}, \theta_i) = 1$, so (3.9) is equal to $\theta_i - \frac{1-\sum_{k=i}^{N}\mu_k}{\mu_i}\Delta\theta$. On histories in which types are not constant, any quantity is optimal. Since these quantities neither affect the surplus nor the rents of the agent they do not enter the objective function, (3.6).[14] The quantity $q^*(\theta_i)$ is equal to the optimal quantity that would be offered in a static model with $T = 1$. This observation was first made by Baron and Besanko [1984]. For future reference, note that this is only one of the possible solutions.

Example 3 (cont.). With two types, (3.9) implies that $q^*(\theta_i | h^{t-1}) = \theta_i$ if $\theta_i = \theta_H$ and/or θ_H is a realization in h^{t-1}. For the remaining history, $\widetilde{h}^{t-1}$, in which the type is always θ_L, we have $q^*\left(\theta_L | \widetilde{h}^{t-1}\right) = \theta_L - \frac{\mu_H}{\mu_L}\left(\frac{F(\theta_L|\theta_L) - F(\theta_L|\theta_H)}{F(\theta_L|\theta_L)}\right)^{t-1} \Delta\theta$. In this case the FO-optimal contract is efficient for all histories except the lowest in which the type is θ_L. Along the lowest history in which quantities are distorted, the distortion is proportional to $\left(\frac{F(\theta_L|\theta_L) - F(\theta_L|\theta_H)}{F(\theta_L|\theta_L)}\right)^{t-1}$, which is less than 1, and so it vanishes as $t \to \infty$.

Example 4 (cont.). In the first period, we have $q^*(\theta_i) = \theta_i - \frac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \Delta\theta$, as in the static model, and in period 2, $q^*(\theta_i | \theta_j) = \theta_i - \frac{1 - \sum_{k=j}^{N} \mu_k}{\mu_j} \frac{F(\theta_i|\theta_j) - F(\theta_i|\theta_{j-1})}{f(\theta_i|\theta_j)} \Delta\theta$.

Some distinct characteristics easily emerge from (3.8) even without assuming that it admits a unique solution. Since the right hand side of (3.8) is non-negative, the contract is always distorted downward, at least weakly: so, analogous to the static case, we never have overprovision, but we can have underprovision. Moreover, the right hand side becomes zero when the type becomes θ_0, the highest type (since $D(h^{t-1}, \theta_i) = 0$). In this case, $s_q(\theta_i, q^*(\theta_i | h^{t-1})) = 0$ and the contract is efficient in all following periods, a phenomenon that has been called "Generalized No-Distortion at the Top". For any other history, the quantities are distorted strictly below the

[14] In the rest of the paper we assume that types have full support so (3.7) is always well defined. With perfect persistence, for histories in which types change, $D(h^t)$ is indeterminate: in this cases both the numerator and the denominator of $D(h^t)$ are zero. These histories occur with zero probability, so the associated quantities are irrelevant.

efficient level. The distortion is exactly equal to $\left[\sum_{k=0}^{j-1} \mu_k/\mu_i\right] D(h^{t-1}, \theta_i)\Delta\theta$: this formula is complicated because the wedge is state contingent and it depends on the entire history.

3.4 When does the first-order approach work?

Given the (relatively) simple characterization of Proposition 1, the imperative question is: when is it without loss of generality to focus on the first-order approach? From previous work we know that there are a number of cases in which the first-order approach works and so the optimal contract coincides with (3.8). Should these cases be seen as the "standard cases", or are they special examples? In the remainder of this section we attempt to answer this question exploring when the first-order approach is valid.

To verify the validity of the FO-approach we need to establish that the solution of (3.6) satisfies the full set of constraints in (3.3). Corollary 1 tells us that the agent's rents are functions only of the quantities, so the set of constraints also depends only on q. We let $C(\mathbf{q})$ denote this set of constraints. The first-order optimal contract is defined by $\mathbf{q}(\Theta, \mu, F)$, function only of the fundamentals. It follows that a necessary and sufficient condition for the validity of the first-order approach is that the set of fundamentals satisfy the family of inequalities defined by $C(\mathbf{q}(\Theta, \mu, F))$.

The key question is whether these conditions define reasonably interesting economic environments for which the FO-approach works. The following result provides a unified framework to interpret existing "possibility results" for the FO-approach. Let $q(h^t) = q\left(h_t^t | h^{t-1}\right)$ be an allocation after history h^t, and let $h^t \succeq \widehat{h}^t$ if $h_j^t \geq \widehat{h}_j^t$ $\forall j \leq t$. We have:

Definition 2. *An allocation is monotonic if* $q(h^t) \geq q(\widehat{h}^t)$ *for any* $h^t \succeq \widehat{h}^t$.

A simple sufficient condition for the validity of the first-order approach can now be stated.[15]

Proposition 2. *The envelope formula (3.5) and monotonicity of the FO-optimal contract are sufficient for implementability.*

Proposition 2 directly parallels the well known results in static environments that show that local incentive compatibility (i.e. the envelope formula) and monotonicity of the allocation are necessary and sufficient for implementability. The result is however weaker for two reasons: first the monotonicity condition is stronger than in a static environment, since it involves all histories following a report; second, the result is only sufficient. The problem with Proposition 2 is that it is useful only to the extent that it is easy to apply. The are a number of applications in which the FO-optimal contract is indeed monotonic.

Example 1 and 2 (cont.). When types are i.i.d., the contract is history independent and monotonic in all periods $t > 1$ (since it coincides with the efficient allocation). The contract is also monotonic in the type at $t = 1$ if the optimal static contract is monotonic: this is always the case if, for example, the prior satisfies the monotone hazard rate condition and $u_{q\theta\theta} \leq 0$, a condition satisfied, for example, by a uniform prior and $u(\theta, q) = \theta q$. When types are constant, the repetition of the first-order optimal static contract is a FO-optimal contract (although it is not unique), which is monotonic if the first-order optimal static contract is monotonic. It follows that under standard assumptions that guarantee monotonicity in θ of the first-order optimal static contract, FO-optimal contract is an optimal dynamic contract both when types

[15] This result generalizes, to an environment with N types, the method used in Battaglini [2005] to establish the sufficiency of (3.5) for $N = 2$. In Battaglini [2005] it is shown that a weaker monotonicity condition than the one in Proposition 2 is actually sufficient. The condition requires that for any history, the marginal of expected utilities are non decreasing in the current type (see Step 1 of Claim 2 in Battaglini [2005]). This condition is implied by the monotonicity of the allocation as defined in Definition 2. Analogous monotonicity results for continuous types are presented by Pavan, Segal and Toikka [2014].

are constant and when they are i.i.d .

Example 3 (cont.). As discussed in the previous section, with two types the FO-optimal contract is efficient in all histories except the history in which types' realizations are all θ_L. This history is also the "lowest history" according to the order $\succeq$. It follows that the contract is monotonic according to Definition 1, and so the FO-optimal contract is optimal.[16]

Example 5: AR(1) model. Besanko [1985] and more recently Pavan, Segal and Toikka [2014] assume an $AR(1)$ model in which $\theta_t = \gamma\theta_{t-1} + \varepsilon_t$, where ε_t is the realization of an i.i.d. random variable and $\gamma \in (0,1)$. The Markovian framework developed above can be easily adapted to generalize this environment to non i.i.d. shocks. Here we present a two period model to drive home the point in a simple fashion. In both periods, the "shocks" have support $\theta_0, ...\theta_N$, with $\theta_k - \theta_{k+1} = \Delta\theta$; in the first period the realization is $\theta_1 = \theta_i$ with prior probability μ_i, in the second period $\varepsilon_2 = \theta_j$ with probability α_{ij} when $\theta_1 = \theta_i$. When $\alpha_{ij} = \alpha_{kj}$ for any i, j, k, we have i.i.d. shocks and the model is equivalent to the models presented in Besanko [1985] and Pavan, Segal and Toikka [2013]; in general, however, it is natural to expect α_{ij} and α_{kj} not to be equal.

When we consider only local incentive constraints, it is easy to show that they must hold as equalities. In period 2, we have $U(\theta_{k,2}|\theta_i) = U(\theta_{k+1,2}|\theta_i) + \Delta\theta q(\theta_{k+1,2}|\theta_i)$, where $U(\theta_{k,2}|\theta_i)$ and $q(\theta_{k,2}|\theta_i)$ are respectively the second period utility and quantity of the agent when the realization at $t = 1$ and $t = 2$ are θ_i and $\theta_{k,2}$.[17] Without loss of generality we can set $U(\theta_{N,2}|\theta_i) = 0$, so that $U(\theta_{k,2}|\theta_i) = \Delta\theta \sum_{l=k+1}^{N} q(\theta_{l,2}|\theta_i)$.

[16]Boleslavsky and Said [2013] generalize this two type model assuming continuous types at $t = 0$ with binary shocks in the following periods. They also consider a version with a continuum of shocks, but in this case they directly assume that the quantities are monotonic in the reported values.

[17]So, with realization θ_i in period 1, we have that $\theta_{k,2} \in \{\gamma\theta_i + \theta_0, .., \gamma\theta_i + .\theta_N\}$.

Similarly, in the first period, we have:

$$U(\theta_i) = U(\theta_{i+1}) + \left[\Delta\theta q(\theta_{i+1}|h^0) + \delta\gamma\Delta\theta \sum_{k=0}^{N} \alpha_{ik} q(\theta_{k,2}|\theta_{i+1})\right] \tag{3.10}$$

$$+\delta \sum_{k=0}^{N} \left(\alpha_{ik} - \alpha_{(i+1)k}\right) U(\theta_{k,2}|\theta_{i+1}) \tag{3.11}$$

The expected utility of the agent with type θ_i is equal to the utility of type θ_{i+1} plus an *informational rent*. The informational rent can be decomposed into two parts. First, we have a deterministic part $\Delta\theta q(\theta_{i+1}|h^0)+\delta\gamma\Delta\theta \sum_{k=0}^{N} \alpha_{ik} q(\theta_{k,2}|\theta_{i+1})$: the realization of type at time 1 affects rents at $t = 1$ (i.e., $\Delta\theta q(\theta_{i+1}|h^0)$), but it effects rents at $t = 2$ as well, in a way that is proportional to γ (i.e., $\delta\gamma\Delta\theta \sum_{k=0}^{N} \alpha_{ik} q(\theta_{k,2}|\theta_{i+1})$). Second, we have the stochastic part $\delta \sum_{k=0}^{N} \left(\alpha_{ik} - \alpha_{(i+1)k}\right) U(\theta_{k,2}|\theta_{i+1})$. This term depends only on the fact that types i and $i+1$ have different expectations on the probability of the shock ε_t at $t = 2$. With i.i.d. shocks the stochastic term of the information rent is zero. The distortions are then *exclusively* deterministic. The first-order optimal quantities are given by $\theta_1 - \Delta\theta \left(1 - \sum_{k=i}^{N} \mu_k\right)/\mu_i$, in period 1, and $\theta_2 - \gamma\Delta\theta \left(1 - \sum_{k=i}^{N} \mu_k\right)/\mu_i$ in period 2 (when θ_i is the realization in the first period). It is immediate to observe that, under a monotone hazard rate assumption on the prior, quantities are monotonic in the sense of Definition 1. It follows from Proposition 2 that the first-order approach works and these quantities describe the optimal contract.

Example 6: AR(k) model and its variations. When the shock ε_t is i.i.d. we can easily generalize the analysis to T periods following the same steps as in Example 5. In this case we can verify that the quantity at time t is

$$\theta_t - \frac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \cdot \gamma^{t-1}\Delta\theta \tag{3.12}$$

when θ_i is the realization in the first period. Distortions, therefore, are history dependent- they depend only on time through γ^{t-1}. Note that θ_t is the first best efficient quantity: the optimal contract is characterized by deterministic distortions that are independent of the Markov process governing the evolution of types. Given this, it is easy to extend the analysis to the k-order autocorrelation case: $\theta_t = \sum_{j=0}^{k} \gamma_j \theta_{t-j} + \varepsilon_t$. The examples can also be extended to a non-linear case in which $\theta_t = l_1(\theta^{\mathbf{t-1}}, \mathbf{q^{t-1}}) + l_2(\theta^{\mathbf{t-1}}, \mathbf{q^{t-1}})\varepsilon_t$, where $l_i(\theta^{\mathbf{t-1}}, \mathbf{q^{t-1}})$ $i = 1, 2$ are both functions of the sequences of types and quantities up to $t-1$. To see this point, note that at time t the terms $\sum_{j=0}^{k} \gamma_j \theta_{t-j}$ or $l_1(\theta^{\mathbf{t-1}}, q^{\mathbf{t-1}})$ are just constants for all types, so they do not have any effect on incentives to reveal the true type and $l_2(\theta^{\mathbf{t-1}}, \mathbf{q^{t-1}})$ disappears since ε_t is an i.i.d. random variable. In all these cases, the key assumption is that the shock is an independent linear addition to the agent's type.[18]

The examples presented above show that the first-order approach can be extended to study quite complex dynamic environments. All the examples, however, can be reconducted to two basic assumptions. The environment studied in Besanko [1985] allows for many possible types (in fact a continuum), but assumes that types change because of linearly additive stochastic shocks uncorrelated with the agent's type. In this environment the shocks are irrelevant for the equilibrium distortions, which are independent of the history of realized types (except for the first).[19] The environment studied in Battaglini [2005] allows the conditional distributions of the types to depend on the type, but limits the analysis to two types only. In this case the optimal contract is history dependent. These two environments have a common feature: in all these cases the FO-optimal allocation is monotonic. In the next section, however, we show that this is not a general property of FO-optimal contracts.

[18]Multiplicative independent shocks also share a similar structure, see Coutry and Li [2000].

[19]We will return on the importance of these assumptions for the first-order approach to work in Section 3.5.1.

3.5 The limits of the first-order approach

In static environments monotonicity only requires that the quantity is non-decreasing in the type. This condition is satisfied under standard regularity conditions. Notably, a sufficient condition for the monotonicity of the optimal contract is that the prior satisfies the monotone hazard rate condition and that $u_{\theta q}$ is non-decreasing in θ. The examples in the previous section may suggest that monotonicity of the optimal contract is a feature of dynamic contracts as well. Dynamic environments, however, are different and monotonicity should not be expected even in the simplest examples.

To see this, consider an example with two periods and Mussa and Rosen [1978] preferences. Lets us assume 3 types, $\theta_0 = \theta_H$, $\theta_1 = \theta_M$ and $\theta_2 = \theta_L$ with $\theta_H > \theta_M > \theta_L$ and transition probabilities $f(\theta_i | \theta_j) = \alpha$ and $f(\theta_i | \theta_j) = (1-\alpha)/2$ for $i \neq j$. These simple transition probabilities satisfy first-order stochastic dominance, so they preserve "order" in the stochastic evolution of types. From Example 4 in Section 3.3 we have:

$$\begin{aligned} q^*(\theta_M | \theta_M) &= \theta_M - \tfrac{1-\sum_{k=L,M}\mu_k}{\mu_M} \frac{F(\theta_M|\theta_M) - F(\theta_M|\theta_H)}{f(\theta_M|\theta_M)} \Delta\theta \\ &= \theta_M - \frac{\mu_H}{\mu_M} \frac{3\alpha-1}{2\alpha} \Delta\theta < \theta_M. \end{aligned} \tag{3.13}$$

Monotonicity would require $q^*(\theta_M | \theta_M) > q^*(\theta_M | \theta_L)$. On the contrary, we have:

$$q^*(\theta_M | \theta_L) = \theta_M - \tfrac{1-\mu_L}{\mu_L} \frac{F(\theta_M|\theta_L) - F(\theta_M|\theta_M)}{f(\theta_M|\theta_L)} \Delta\theta = \theta_M,$$

since $F(\theta_M|\theta_M) = F(\theta_M|\theta_L) = \frac{3\alpha-1}{2}$. So the FO-optimal contract is not monotonic with respect to the realization at $t = 1$.

To understand why the seller finds it optimal to offer a non-monotonic contract, consider the role of distortions in screening problems. Starting from the surplus maximizing contract, consider the effect a marginal reduction in $q^*(\theta_M|\theta_M)$ (see the left panel of Figure 3.1). The principal distorts $q^*(\theta_M|\theta_M)$ only if the change induces a reduction in the agent's rents at $t = 1$ since the change implies a costly reduction in surplus. The reduction in $q^*(\theta_M|\theta_M)$ reduces $U(\theta_H|\theta_M)$ at $t = 2$ (see the left panel of Figure 3.1).[20] If this reduction is, say Δ, then the effect on rents at $t = 1$ is a reduction in $U(\theta_H)$ by $\delta[f(\theta_H|\theta_H) - f(\theta_H|\theta_M)]\Delta$, that is $\delta(3\alpha - 1)\Delta/2$. This follows from the fact that by reducing $U(\theta_H|\theta_M)$, the seller makes a deviation to θ_M less profitable for a type θ_H at $t = 1$.[21] The seller chooses the distortion to solve the trade-off between the marginal benefit of reducing the rent of θ_H, and the marginal cost in terms of surplus reduction. Indeed, the first-order condition characterizing $q^*(\theta_M|\theta_M)$ is:

$$\delta \cdot \mu_M \cdot \alpha \cdot [\theta_M - q^*(\theta_M|\theta_M)] = \delta \cdot \mu_H \cdot \frac{3\alpha - 1}{2} \tag{3.14}$$

Since $\theta_M - q^*(\theta_M|\theta_M)$ is the derivative of the surplus generated in period 2, the left hand side of (3.14) is the expected marginal reduction in surplus due to reduction in $q^*(\theta_M|\theta_M)$; the right hand side of (3.14) is the expected reduction in paid rents.[22]

Consider now a reduction in $q^*(\theta_M|\theta_L)$ that keeps the rent of the θ_L type constant at $t = 1$ (see the right panel of Figure 3.1). By the same logic as above, a marginal reduction in $q^*(\theta_M|\theta_L)$ induces a decrease in the rents of θ_M at $t = 1$ equal to

[20] From the binding IC_{HM} constraint, at t=2 we have $U(\theta_H|\theta_M) = U(\theta_M|\theta_M) + \Delta\theta \cdot q(\theta_M|\theta_M)$.

[21] If the reduction in $U(\theta_H|\theta_M)$ is Δ, then the expected reduction in the rent of type θ_M at $t = 1$ is $\delta f(\theta_H|\theta_M)\Delta$. The principal, however, can not allow a reduction in $U(\theta_H|\theta_M)$ without other changes in the contract. Incentive compatibility at $t = 1$ requires that the rent of type θ_M is equal to the rent this type would receive if he reported θ_L: to preserve incentive compatibility $U(\theta_M)$ must remain constant. To compensate for the expected reduction in the rent, the principal reduces the price paid by θ_M at $t = 1$ by $\delta f(\theta_H|\theta_M)\Delta$. The benefit of reducing $q^*(\theta_M|\theta_M)$, therefore, is captured exclusively by the fact that the outside option at $t = 1$ for θ_H is reduced. After the change, the rent of θ_H at $t = 1$ (that in equilibrium is equal to the utility of falsely reporting to be θ_M) is therefore reduced by a net $\delta[f(\theta_H|\theta_H) - f(\theta_H|\theta_M)]\Delta$.

[22] With respect to the right hand side of (3.14), note that $\mu_M \cdot \alpha$ is the probability of history $h^2 = \{\theta_M, \theta_M\}$. With respect to the left hand side, note that the expected reduction in paid rents is the change in $U(\theta_H)$ times the probability of a type θ_H at $t = 1$.

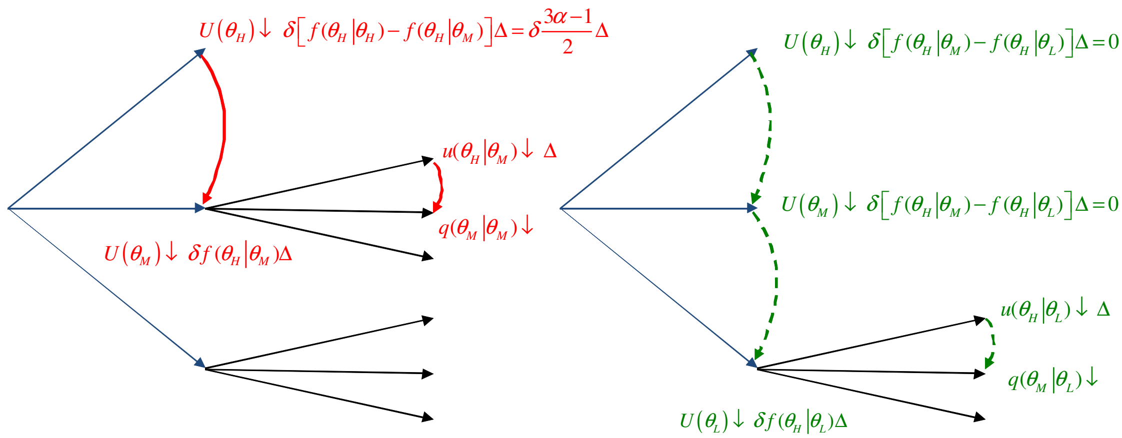

Figure 3.1: The simple economics behind optimal non-monotonic allocations.

$\delta\left[f\left(\theta_H\left|\theta_M\right.\right)-f\left(\theta_H\left|\theta_L\right.\right)\right]\Delta$. The rent of θ_H is a linear function of $U(\theta_M)$, so it is also reduced by the same amount. In our model, however we have that $f\left(\theta_H\left|\theta_M\right.\right)$ and $f\left(\theta_H\left|\theta_L\right.\right)$ are both equal to $(1-\alpha)/2$, so the net benefit of the reduction in $q^*\left(\theta_M|\theta_L\right)$ is zero. It follows that the equilibrium condition is:

$$\delta\cdot\frac{1-\alpha}{2}\cdot\left[\theta_M-q^*\left(\theta_M\left|\theta_L\right.\right)\right]=0 \tag{3.15}$$

where $\delta\cdot\frac{1-\alpha}{2}\cdot\left[\theta_M-q^*\left(\theta_M\left|\theta_L\right.\right)\right]$ is the expected marginal reduction in surplus due to the reduction in $q^*\left(\theta_M|\theta_L\right)$.

Comparing (3.14) to (3.15) we can see that $q^*\left(\theta_M|\theta_L\right)>q^*\left(\theta_M|\theta_M\right)$ because the the marginal effect on rents of distorting $q^*\left(\theta_M\left|\theta_L\right.\right)$ is smaller than the effect of $q^*\left(\theta_M|\theta_L\right)$ for $\alpha>1/3$. A reduction in $q^*\left(\theta_M|\theta_L\right)$ induces a reduction in the rent of a type θ_H at $t=2$: a type θ_M is however as likely likely to become θ_H as a type θ_L (since $f\left(\theta_H\left|\theta_M\right.\right)=f\left(\theta_H\left|\theta_L\right.\right)$), so this has no screening effect on the rent extracted by θ_M at $t=1$. (In general, $f\left(\theta_H\left|\theta_M\right.\right)-f\left(\theta_H\left|\theta_L\right.\right)$ is always small when types are highly persistent, since in this case both $f\left(\theta_H\left|\theta_M\right.\right)$ and $f\left(\theta_H\left|\theta_L\right.\right)$ are both close to zero). Similarly, a reduction in $q^*\left(\theta_M|\theta_M\right)$ induces a reduction in the rent of a type θ_H in $t=2$: for $\alpha>1/3$, however, θ_H is more likely to remain θ_H than θ_M is to become θ_H; this makes a distortion on $q^*(\theta_M|\theta_M)$ a more effective screening device than a distortion on $q^*(\theta_M|\theta_L)$.

In the example presented above we have assumed a particular transition function $f_\alpha(\theta_j|\theta_i)$ in which the probability of persistence is the same for all types ($f_\alpha(\theta_i|\theta_i) = \alpha$) and all deviations are equally likely ($f_\alpha(\theta_j|\theta_i) = \frac{1-\alpha}{N}$ for $i \neq j$). Does the phenomenon illustrated by the example extend to general transition functions? To address this question consider the general set Λ of all possible transition functions $f_\alpha(\theta_j|\theta_i)$ satisfying Assumption 2 and parametrized by $\alpha \in [0,1]$ such that for all i, $f_\alpha(\theta_i|\theta_i) \to 1$ as $\alpha \to 1$.[23]

Definition 3. *We say that a property holds for a* generic *transition function in* Λ *if it holds for an open and dense set of functions in* Λ.

This is the standard definition of genericity in this environment.[24] Our first result proves that, for a generic transition function, the optimal contract is non-monotonic when types are highly persistent.

Proposition 3. *For any* μ, δ, $|\Theta| > 2$, $T > 2$, *and a generic transition function in* Λ, *there exists an* $\alpha^* < 1$ *such that the FO-optimal contract is not monotonic for any* $\alpha > \alpha^*$.

To grasp the intuition behind this result, consider the FO-optimal contract when $u(\theta, q) = \theta q$. From (3.9), in an interior solution, we have:

$$q^*\left(\theta_i|h^{t-1}\right) = \theta_i - \frac{1-\sum_{k=j}^{N}\mu_k}{\mu_j} D(h^{t-1}, \theta_i)\Delta\theta$$

[23]Note that since we impose no additional restrictions, the probabilities of persistence of different types may be different for $\alpha < 1$ and they can even converge to one at different speeds: α is just an index of the level of persistence of the stochastic process.

[24]Endowed with a sup norm, the space of transition functions Λ is a complete metric space. The complement of an open and dense set in Λ is a set of first category. The Baire Category Theorem guarantees that these sets have empty interior and therefore are topologically small (Royden [1988], ch.7.8).

Intuitively (and as formally proven in the proof of Proposition 3), for a generic transition probability function the terms $D(h^{t-1}, \theta_i)$ are history dependent, even as persistence converges to one. In particular, if $h^{t-1}(\theta_i)$ is the history in which all realizations are θ_i, then there is always an h^{t-1} with $h_1^{t-1} = \theta_i$, $h^{t-1} \prec h^{t-1}(\theta_i)$ and either $D(h^{t-1}, \theta_i) < D(h^{t-1}(\theta_i), \theta_i)$ or $D(h^{t-1}, \theta_i) > D(h^{t-1}(\theta_i), \theta_i)$. In the first case it is clear that the FO-optimal contract is non-monotonic and the result is proven. The key step in the proof of Proposition 3 is to show that if $h^{t-1} \preceq h^{t-1}(\theta_i)$ and $D(h^{t-1}, \theta_i) > D(h^{t-1}(\theta_i), \theta_i)$, then there must be a history $\underline{h}^{t-1}$ with $\underline{h}_1^{t-1} = \theta_{i+1}$, $\underline{h}^{t-1} \succ h^{t-1}(\theta_{i+1})$ and $D(\underline{h}^{t-1}, \theta_{i+1}) > D(h^{t-1}(\theta_{i+1}), \theta_{i+1})$ when types are sufficiently persistent. This implies $q^*\left(\theta_{i+1}|h^{t-1}(\theta_{i+1})\right) > q^*\left(\theta_{i+1}|\underline{h}^{t-1}\right)$ and, alas, a failure of monotonicity. The reason why the result holds generically is that the concept of monotonicity required is so strong that even a arbitrarily small difference between $D(h^{t-1}, \theta_i)$ and $D(h^{t-1}(\theta_i), \theta_i)$ is sufficient to make it fail, independently of its sign.

Does Proposition 3 imply that the FO-approach generically fails when types are highly persistent? It is easy to see that a failure of monotonicity can not alone be sufficient for the first-order approach to fail. When δ is small, the future becomes irrelevant and the problem is essentially static. What happens when the future is sufficiently important? In the next result we use Proposition 3 to show that when types are highly persistent and the future is sufficiently important, then even a small failure of monotonicity is sufficient to make the FO-approach invalid. We have:

Proposition 4. *For any μ, $|\Theta| > 2$ and a generic transition probability function, there exists an $\alpha^* < 1$, $T^* > 2$, and $\delta^* < 1$ such that the first-order approach fails to be verified for any $\alpha > \alpha^*$, $T \geq T^*$ and $\delta > \delta^*$.*

The proposition shows that with highly persistent types and a long horizon, the effect of even a small failure of monotonicity is highly magnified as $\delta \to 1$, to the point of inducing an effect on the expected utility that is significant enough to violate a global incentive constraint. The fact that the first-order approach fails only if types

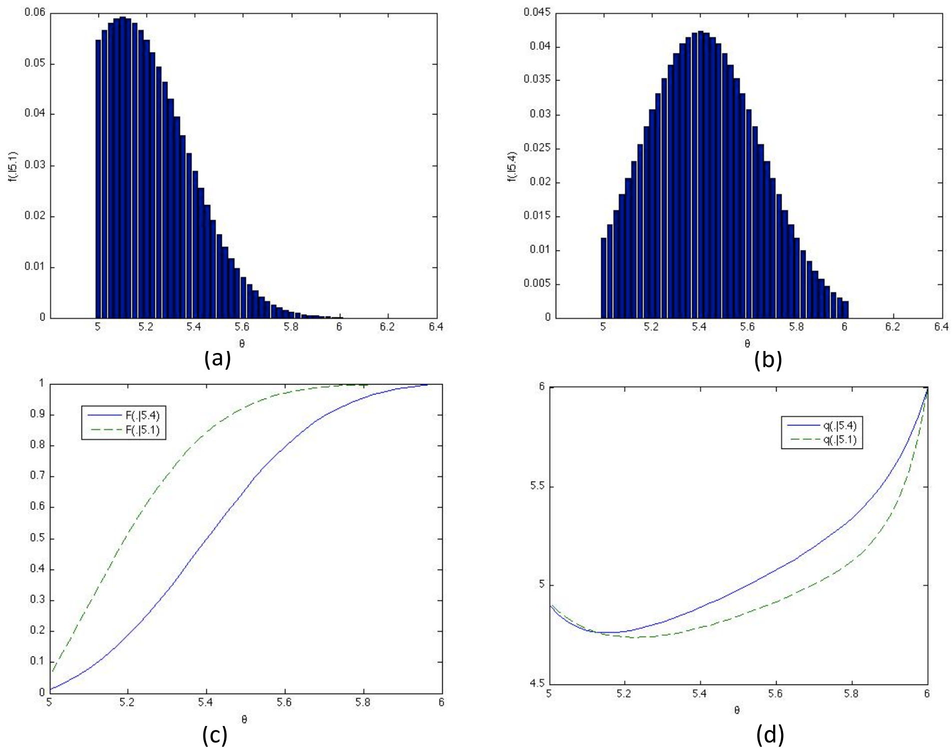

Figure 3.2: Example 7, the truncated normal case. $\sigma = 0.25$, $\Delta\theta = 0.025$.

are sufficiently persistent and expected payoffs are sufficiently important should not be surprising. As we have seen in Example 1, the first-order approach always works when types are sufficiently serially uncorrelated: in this case types have small private information about the future, so there is no point in imposing distortions on future quantities. Similarly, if agents are impatient and the time horizon is short, the model is close to being static.

Interestingly, it is easy to compute examples in which very limited serial correlation is sufficient to induce a failure of the first-order approach. To illustrate this point, we conclude this section with three natural examples.

Examples 7. Assume that the type in the first period, θ_1, is uniformly distributed on $\Theta = [5, 6]$ and the distribution in the second period is a (truncated) normal $f_\alpha(\theta_2 | \theta_1) = \frac{A(\theta_1)}{\sigma}\Phi(\frac{\theta_2 - \theta_1}{\sigma})\Delta\theta$ where Φ is a standard normal density with variance σ and $A(\theta_1)$ is chosen so that the distribution assigns probability one on

Θ.[25] (The specific values chosen for the support are obviously irrelevant and chosen only as examples). In this case, the probability that a type remains constant is $f_\alpha(\theta_i|\theta_i) = \frac{A(\theta_i)}{\sigma}\Delta\theta\Phi(0)$, a function of the first period realization. It is easy to verify the probability of persistence converges to one as $\sigma \to 0$ (in the notation used above, the process can therefore be parametrized by $\alpha = 1-\sigma$). The top panels of Figure 3.2 illustrate $f_\alpha(\theta_j|\theta_i)$ for two values: $\theta_i = 5.1$ (top left panel) and $\theta_i = 5.4$ (top right panel) The bottom right and left panel of the figure shows the FO-optimal contracts at $t = 2$ after histories $\theta_i = 5.1$ and $\theta_i = 5.4$, respectively. The contract is not monotonic: it is not monotonic with respect to the realization at $t = 2$ (this can be seen from the fact that the lines are not non-decreasing); and it is not monotonic in the realization at $t = 1$ (this can be seen form the fact that the contracts intersect at $t = 2$). It can also be verified that the FO-optimal contract is not incentive compatible.

Examples 8 and 9. Assume that the type in the first period, θ_1, is uniformly distributed and consider now the transition probabilities $f_\alpha(\theta_j|\theta_i) = \alpha e^{-\frac{(\theta_j-\theta_i)^2}{\sigma_i(\alpha\Delta\theta)}}\Delta\theta$, and $f_\alpha(\theta_j|\theta_i) = \frac{\alpha\Delta\theta}{1+\sigma_i|\theta_j-\theta_i|}$, where σ_i chosen so that the probabilities sum to one. In this case, $f_\alpha(\theta_i|\theta_i) = \alpha\Delta\theta$ so the probability of persistence is identical for all types; it is the variance of the distribution that is adjusted so that f_α assigns probability one on Θ_2. As it is straightforward to verify, we have $\sigma_i \to 0$ as $\alpha\Delta\theta \to 1$. Figures 3.3 and 3.4 illustrate $f_\alpha(\theta_j|\theta_i)$ for two values: $\theta_i = 5$ (top left panel) and $\theta_i = 5.1$ (top right panel) and compares the two implied distribution functions (bottom left panel) assuming $\Theta_1 = [5,6]$ and $\Theta_2 = [4.5, 6.5]$. The bottom right panel of the figures illustrates the FO-optimal contracts at $t = 2$ after histories $\theta_i = 5$ and $\theta_i = 5.1$, respectively. As in Example 7, the contract is not monotonic and the FO-optimal

[25] To obtain a discrete density that can be applied to any size $\Delta\theta$ (even arbitrarily small), we discretize a continous density $f(\theta)$. As standard, in this case the probability of type θ_i is equal to $f(\theta_i)\Delta\theta$, i.e. the "histogram" approximation of the continuous density.

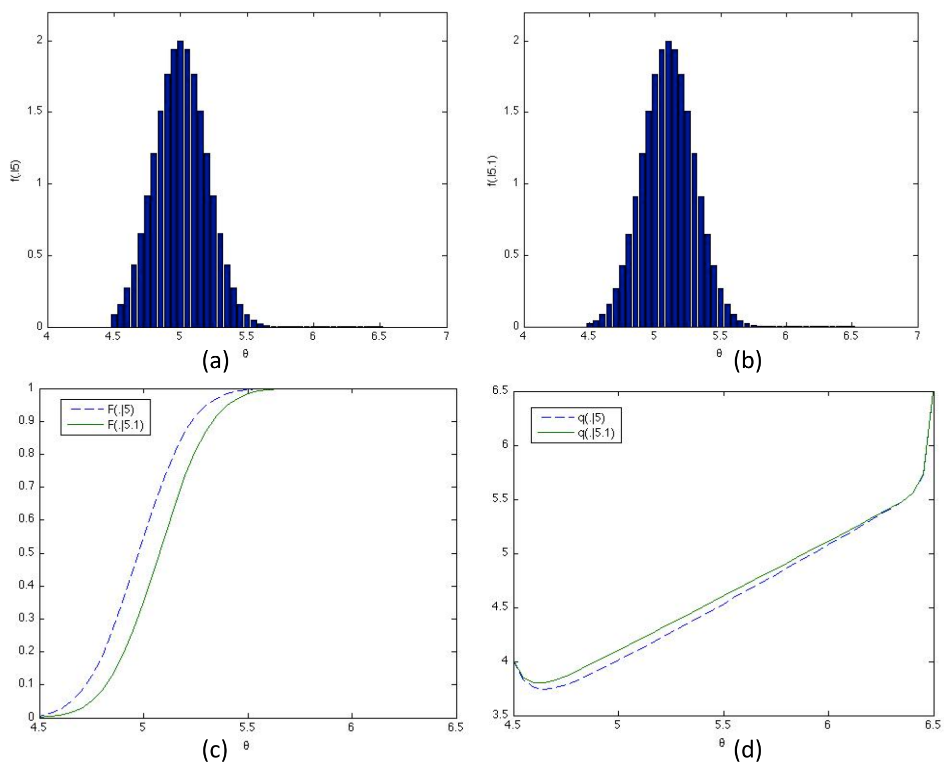

Figure 3.3: Example 8, the exponential case. $\alpha = 2$, $\Delta\theta = 0.05$.

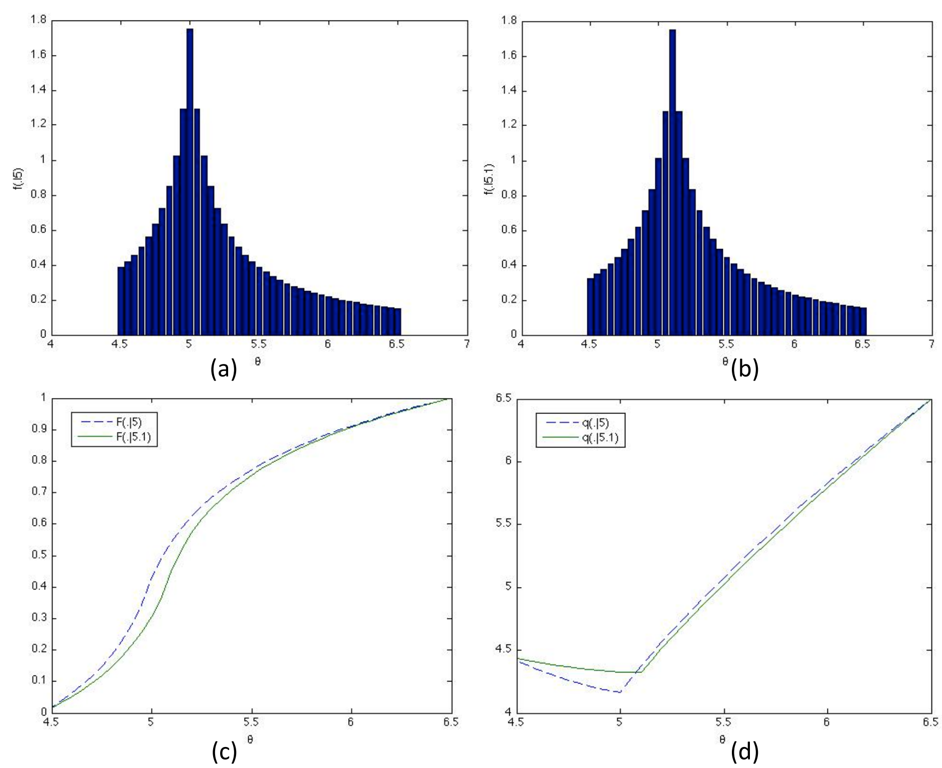

Figure 3.4: Example 9, the hyperbolic case. $\alpha = 1.75$, $\Delta\theta = 0.05$.

contract associated to this case is not incentive compatible.

It is interesting to note that both in Examples 7, 8 and 9 the contract is not monotonic despite the fact that the transition probabilities have very little persistence. Finally, the discussion above, Propositions 3 and 4, and non-monotonicity in the examples are all valid independent of the prior μ, re-affirming our assertion that the failure of the first-order approach is not a technical irregularity, but a consequence of the added structure that dynamics present to the economic problem of contracting.

3.5.1 Discussion

We conclude this section with a few remarks on Propositions 3 and 4.

Perfectly persistent shocks. As we have seen in the previous sections, the first-order approach *always* works when types are perfectly persistent; Proposition 4, however, shows that the FO-approach does not generically work when types are highly persistent. How is this possible? The key to understanding this apparent contradiction is to realize that when types are constant, the repetition of the optimal static contract is only one of the many possible solutions: in histories which occur with *exactly* zero probability, the quantities are irrelevant and so they can be set to any arbitrary number, for example, equal to the static optimum. On the contrary, when types are highly persistent, but probabilities off the main diagonal are not exactly zero, quantities can not be set arbitrarily in these histories. The effect of these histories on the agent's rents is small, but so is the effect of these quantities on the surplus. Typically, the quantities are uniquely defined along all histories. As persistence converges to one, these quantities along the non-constant histories converge to values that are different from the static optimum and that are non-monotonic. In the example presented at the beginning of Section 3.5, $q_M(M) = \theta_M - \frac{\mu_H}{\mu_M}\frac{3\alpha-1}{2\alpha}$ that converges to $\theta_M - \frac{\mu_H}{\mu_M}$ in the limit; $q_M(L)$, on the contrary, is equal to θ_M for any α: in the limit, therefore, $q_M(M) < q_M(L)$. The problem is that there is a lack of lowerhemicontinuity at the limit with constant types, and some of the limit solutions (including the repetition of the static optimum) can not be seen as the limit of solutions as persistence converges to one.

AR(k) models. To see why monotonicity is a fragile property in $AR(k)$ models consider (3.12). In this formula the terms $D(h^{t-1}, \theta_i)$ are all identically equal to γ^{t-1} and independent of h^{t-1}: trivially, therefore, we have $q^*(\theta_i|h^{t-1}) = q^*\left(\theta_i|\widehat{h}^{t-1}\right)$ for *any* two histories with $h^{t-1} \succeq \widehat{h}^{t-1}$. This however is not a generic or even a plausible property: it follows from the fact that the shocks ε_t are assumed to be i.i.d and linearly additive. If we assume that the distribution of ε_t depends on the past realization, even if the effect of the past realization is very small, then $D(h^{t-1}, \theta_i)$ is history dependent and it is no longer the case that $q^*(\theta_i|h^{t-1}) = q^*\left(\theta_i|\widehat{h}^{t-1}\right)$.

In addition to this, even assuming a constant γ, to make the $AR(k)$ model conceivable we need to assume that the support shifts with the type as in Examples 4 and 5 or alternatively that the type support is unbounded above and below. If we assume a given constant and bounded support then a perfect horizontal translation of the distribution is obviously impossible. Again, it is not generally true that $q^*(\theta_i|h^{t-1}) = q^*\left(\theta_i|\widehat{h}^{t-1}\right)$ for *any* two histories with $h^{t-1} \succeq \widehat{h}^{t-1}$. Example 7 can be seen as an $AR(k)$ model with bounded support in which the shock follows a truncated normal. As evident from Figure 3.2, neither monotonicity nor the FO-approach works in this case.

On serially independent shocks. Pavan Segal and Toikka [2013] suggest that the *AR(k)* model can be seen as an example of a more general class of environments for which the first-order approach works. Their suggestion is based on an original observation by Eso and Szentes [2007] who note that any model with correlated and continuous types can be transformed into an equivalent model with i.i.d. shocks. To see this, note that if the cumulative distribution is $F(\theta_t|\theta_{t-1})$, then assuming that the agent observes θ_t is equivalent to assuming that he or she observes the variable $v_t = F(\theta_t|\theta_{t-1})$ (since $F(\theta_t|\theta_{t-1})$ is increasing and invertible in θ_t): as well known, v_t is a random variable with a uniform distribution on $[0,1]$. Eso and Szentes' [2007] observation is insightful in interpreting the screening contract and useful to derive the envelope formula (see Eso and Szentes [2013]). Unfortunately, however, it does not generalize the insights from the *AR(1)* models, and it does not help to solve the problems of the FO-approach elucidated above. It is useful to illustrate why transforming the stochastic process to an i.i.d. shock does not make the problem more tractable. Assume the utility is $u(\theta, q) = \theta q$. To make the equivalent transformation, we need to substitute $\theta_t = F^{-1}(v_t; \theta_{t-1})$, so we have: $u(v_1, q_1) = \mathbf{F}^{-1}(\mathbf{v}_1) \cdot q_1$, $u(\mathbf{v}^2, q_2) = F^{-1}(v_2; F^{-1}(v_1)) \cdot q_2$ and iterating:

$$u(\mathbf{v}^t, q_t) = F^{-1}(v_t; F^{-1}(v_{t-1}; F^{-1}(v_{t-2}; [...]))) \cdot q_t \tag{3.16}$$

where $\mathbf{v}^t = (v_1, ..., v_t)$. It is clear from (3.16) that, even starting from the simplest utility function, the per period utility of the equivalent transformation is a very complicated, time inseparable function of the entire history of the shocks $\mathbf{v}^t$. The change of variables, from θ_t to v_t, allows one to get rid of serial correlation in the types; the correlation however, does not disappear: it must be incorporated in a transformed utility function. All the problems that induce a failure of the first-order approach in the original problem are just shifted from the distribution function to the transformed per period utility function. The benefit of having independent shocks is compensated by the complications of having these per period utilities.

From discrete to continuous types. While we have focused the analysis on the case with discrete types, there is a strict connection between models with discrete and continuos types; and the same issues discussed above arise in continuous type models as well. Consider a continuous types model with type set $\Theta = \left[\underline{\theta}, \overline{\theta}\right] \subset \mathbb{R}^+$, prior distribution $\Gamma(\theta)$ and transition distribution $F(\theta' | \theta)$. We can define an associated discrete model by defining the type space as $\Theta^N = \{\theta_0, ..., \theta_N\}$ with $\theta_0 = \overline{\theta}$, $\theta_N = \underline{\theta}$ and $\theta_i = \theta_{i+1} + \Delta\theta_N$, the prior as $\Gamma^N(\theta_i) = \Gamma(\theta_i)$ and the transition matrix as $F^N(\theta_j|\theta_i) = F(\theta_j|\theta_i)$. In the online appendix we show that the envelope formula and the FO-optimal contracts of the continuous model can be obtained as limits of the discrete formulas (3.4) and (3.9). We also present a number of solved continuous type examples (including continuous type versions of Examples 8 and 9) to illustrate the problems with the FO-approach.

3.6 What does the optimal contract look like when the first-order approach is invalid?

As we have seen in Section 3.5, even with two periods and three types the FO-optimal contract fails to be monotonic and the FO-approach can not be generally applied. In this section we fully characterize the optimal contract in the motivating example of Section 3.4, in which $f(\theta|\theta) = \alpha$ and $f(\theta|\theta') = \frac{1-\alpha}{2}$ for any $\theta, \theta' \in \{\theta_H, \theta_M, \theta_L\}, \theta \neq \theta'$ and $\alpha > 1/3$. The goal of this section is twofold- to elucidate the structure of optimal contracts that is otherwise elusive in models where the FO-approach can be applied, and to illustrate the trade-offs between rent and efficiency in a dynamic model.

To characterize the optimal contract we focus on a *weakly relaxed program* that constitutes problem (3.3) with $|\Theta| = 3$ and $T = 2$, with the following subset of constraints:

$$IR_L, IC_{HM}, IC_{ML}, IC_{HL}, \tag{3.17}$$
$$IC_{HM}(M), IC_{ML}(M), IC_{LM}(M), IC_{HM}(L), IC_{ML}(L), IC_{LM}(L)$$

where IR_L is the individual rationality constraint of type L at $t = 0$, $IC_{i,j}$ is incentive compatibility constraint requiring that type i doesn't want to misreport being a type

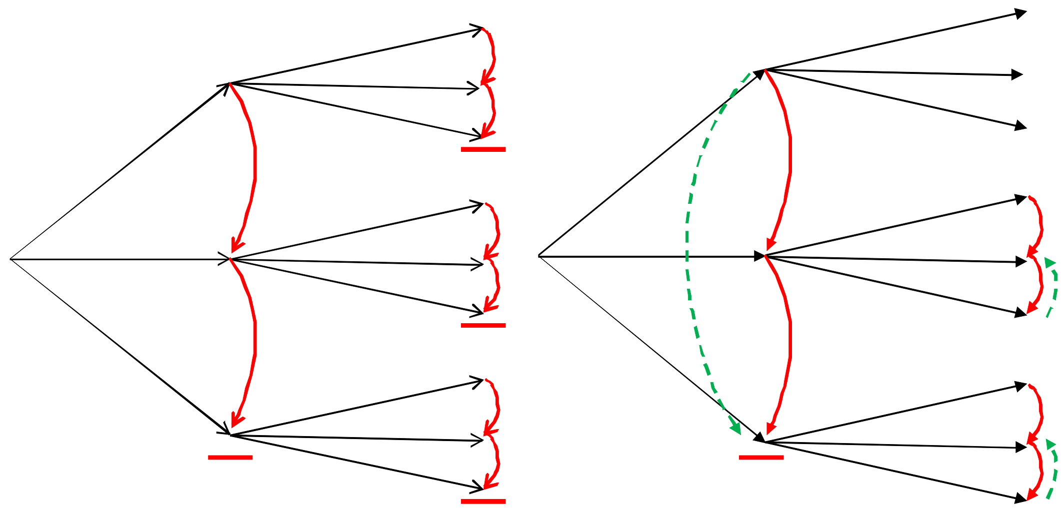

Figure 3.5: The dashed arrows are the constraints in the *WR-program* that are ignored in the Þrst-order approach.

j in period 1, and $IC_{i,j}(k)$ is the incentive compatibility constraint requiring that type i doesn't want to misreport being a type j in period 2, after the agent reports to be a type k in period 1. In contrast to the FO-approach, this problem has two key differences. First, now we are ignoring all the individual rationality constraints of the lowest type in period 2 and incentive compatibility constraints after history H. Second, and most importantly, we are adding three new constraints: the global downward constraint IC_{HL}, and the local upward constraints $IC_{LM}(M), IC_{LM}(L)$ in period 2. The constraint set of the problem is illustrated in the relevant history tree in Figure 3.5. In the following we will refer to this program as the *WR-program*.

Since this is a three type and two period model we simplify notation. Let U_i be the expected utility of type i in the first period and $u_i(h)$ be the expected utility of type i after history h in the second period. Note that since the second period is the terminal period, the expected utility and stage utility are the same. Similarly, we define q_i and $q_i(h)$ to be the first and second period allocations respectively. The following lemma allows to simplify the constraint set:[26]

[26]When only the usual local downward incentive compatibility constraints are

Lemma 2. *In the WR-program, constraints IR_L, IC_{HM}, IC_{ML} bind at the optimum.*

We can now use the equalities implied by Lemma 2 to reduce the number of free variables in the optimization problem. In particular we can eliminate the period 1 utility vectors. Define $\omega_{HM}(i) = u_H(i) - u_M(i)$ and $\omega_{ML}(i) = u_M(i) - u_L(i)$ for $i = M, L$. The variable $\omega_{kl}(i)$ is the net utility of reporting to be type k rather than a type l after history i. Using this notation, we can rewrite the *WR-program* as a maximization problem in which the control variables are the quantities $\mathbf{q}$ and second period marginal utilities ω:

$$\max_{\langle \omega, \mathbf{q} \rangle} \left\{ \begin{array}{c} \sum_{i=H,M,L} \mu_i \left[\theta_i q_i - \frac{1}{2} q_i^2 + \delta \sum_{k=H,M,L} \mu\left(k \,|i\right) \left(\theta_k q_k(i) - \frac{1}{2} q_k(i)^2\right) \right] \\ -\mu_H \left[\Delta\theta q_M + \delta \frac{3\alpha-1}{2} \omega_{HM}(M) \right] \\ -(\mu_H + \mu_M) \left[\Delta\theta q_L + \delta \frac{3\alpha-1}{2} \omega_{ML}(L) \right] \end{array} \right\} \tag{3.18}$$

subject to

$$[\lambda] : \quad \Delta\theta q_M + \delta \frac{3\alpha - 1}{2} \omega_{HM}(M) \geq \Delta\theta q_L + \delta \frac{3\alpha - 1}{2} \omega_{HM}(L)$$

$$[\lambda_{HM}(M)] : \quad \omega_{HM}(M) \geq \Delta\theta q_M(M) \quad | \quad [\lambda_{HM}(L)] : \quad \omega_{HM}(L) \geq \Delta\theta q_M(L)$$

$$[\lambda_{ML}(M)] : \quad \omega_{ML}(M) \geq \Delta\theta q_L(M) \quad | \quad [\lambda_{ML}(L)] : \quad \omega_{ML}(L) \geq \Delta\theta q_L(L)$$

$$[\lambda_{LM}(M)] : \quad \omega_{ML}(M) \leq \Delta\theta q_M(M) \quad | \quad [\lambda_{LM}(L)] : \quad \omega_{ML}(L) \leq \Delta\theta q_M(L)$$

where the variables in the square brackets on the left are the Lagrange multipliers associated with the constraints. Program (3.18) is a standard maximization problem, but it is complicated by a still significantly large number of constraints. The difference between (3.18) and the problem of the first-order approach (3.6) is the global

considered, the following result is immediate. If, for example IC_{HM} were not binding, the principal could simply raise the price that type θ_H is paying. In the *WR-program* the proof of the result is complicated by the additional constraints: reducing type θ_H's rent at $t = 0$ may conflict with IC_{HL}.

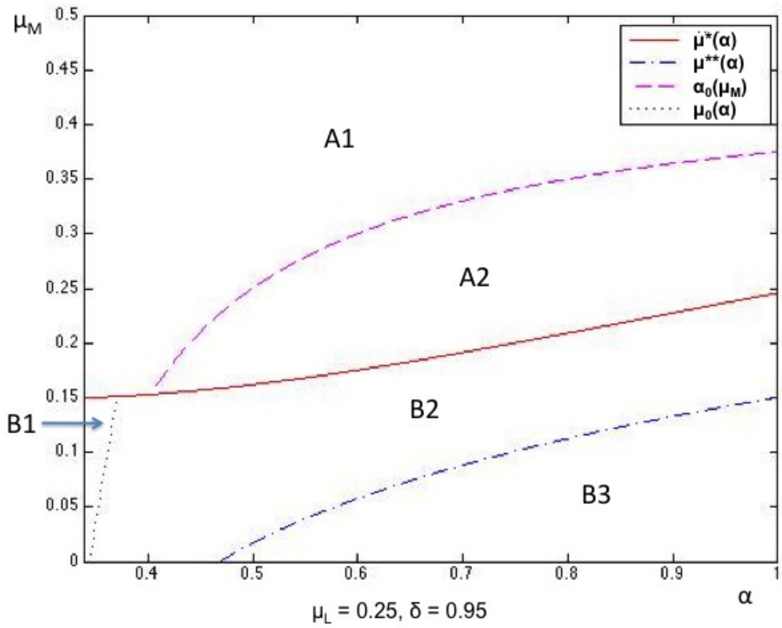

Figure 3.6: Fully characterized contract.

constraint IC_{HL} and the presence of the local upward constraints $IC_{LM}(M)$ and $IC_{LM}(L)$. The latter are essentially *monotonicity conditions* requiring $q_M(h) \geq q_L(h)$ for $h = M, L$.[27] We cannot ignore any of these three constraints. Moreover now we cannot assume without loss of generality that all local downward incentive constraints are binding at $t = 2$: so the envelope formula (3.4) cannot be directly applied. Hence, we still have utilities in the objective function. The next lemma validates our focus on problem (3.18) :

Lemma 3. *A contract is optimal if and only if it solves the WR-program.*

The analysis can be divided into two cases: first the case in which the global constraint can be ignored and so it is sufficient to look at local constraints, i.e. $\lambda = 0$; second, the case in which the global constraint is binding, i.e. $\lambda > 0$.

[27]To see this note that given $IC_{ML}(h)$, $q_M(h) \geq q_L(h)$ if and only if $IC_{LM}(h)$ is satisfied.

3.6.1 Case 1: Local IC is sufficient

The following result characterizes the necessary and sufficient condition for $\lambda = 0$. For a given μ_L and δ, the environment is fully described by two parameters, μ_M, α, and therefore it can be represented in the two dimensional box $(\mu_M, \alpha) \in E(\mu_L) = (0, 1-\mu_L) \times (1/3, 1)$.[28] In the rest of the analysis we will fix μ_L and δ and study how the equilibrium changes as we change μ_M, α. This approach is without loss of generality and it allows for simpler statements (and a graphical representation) of the relevant cases. We have:

Lemma 4. *There exists a threshold $\mu^*(\alpha)$ such that the global incentive constraint IC_{HL} can be ignored if and only if $\mu_M \geq \mu^*(\alpha)$.*

Within the two regions defined by $\mu^*(\alpha)$, the particular shape of the optimal contract depends on the remaining set of binding constraints. Explicit solutions of the optimal quantities for all feasible parameters are presented in Table 1 in the appendix. The following proposition describes what the optimal contract looks like for $\mu_M \geq \mu^*(\alpha)$, when the global constraint can be ignored:

Proposition 6. *Assume $\mu_M \geq \mu^*(\alpha)$. There is a threshold $\alpha_0(\mu_M)$, such that:*

- **Case A1.** *If $\alpha < \alpha_0(\mu_M)$, the optimal contract is fully separating and first-order optimal.*

- **Case A2.** *If $\alpha \geq \alpha_0(\mu_M)$, the optimal contract is fully separating after all histories except M. After this history types M and L are pooled: $q_M(M) = q_L(M)$.*

Regions *A1* and *A2* are illustrated in Figure 3.6 in a simple parametric example, where the threshold $\alpha_0(\mu_M)$ is represented by a dashed line.[29] In region $A1$, the

[28]The thresholds defined below do not depend on the types θ.

[29]In Figure 3.6 we assume $\mu_L = 0.25$ and $\delta = 0.95$.

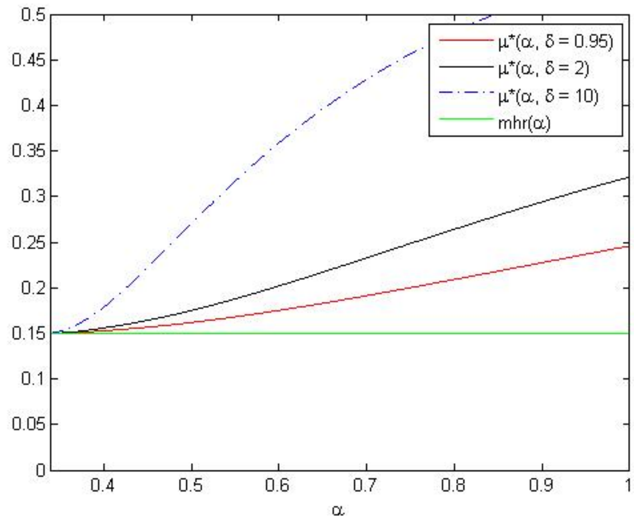

Figure 3.7: Fully characterized contract.

envelope formula is sufficient to characterize the optimal contract. In this case the FO-optimal contract is not monotonic (as in Definition 2), but this lack of monotonicity is not sufficient to cause a failure of incentive compatibility. The contract is not monotonic because $q_M(M) < q_M(L)$. However, given any h, $q_\theta(h)$ is monotonic in θ. In region A2, even though the global constraints can be ignored, the envelope formula is not sufficient to determine the contract since at $t = 2$ we have pooling after history M. It is interesting to note that although pooling makes sure that $q_\theta(h)$ is monotonic in h for all θ, the optimal contract remains non-monotonic with respect to the realization at $t = 1$ (since $q_M(M) < q_M(L)$ in A2 as well).

Figure 3.7 illustrates what happens when we make the payoffs in the second period more important by increasing δ. We know from Proposition 4 that as the future becomes more important, the first-order approach is never valid for high levels of α. A similar phenomenon occurs here: as δ increases, $\mu^*(\alpha)$ shifts up and the region in which local constraints are sufficient shrinks. These higher values of δ should be seen as representative of dynamic models with longer time horizons.

3.6.2 Case 2: Local IC is not sufficient

When $\mu_M < \mu^*(\alpha)$ *both* the global constraint IC_{HL} *and* the local constraints IC_{HM} and IC_{ML} are simultaneously binding in the first period. There are three relevant cases. The following result characterizes the optimal contract in these situations:

Proposition 7. *There exists a threshold $\mu^{**}(\alpha)$ such that:*

- **Case B1&B2**. *Assume $\mu_M \in [\mu^{**}(\alpha), \mu^*(\alpha))$. The optimal contract is fully separating at $t = 1$. There exists a threshold $\mu_0(\alpha)$, such that the optimal contract is fully separating at $t = 2$ as well if $\mu_M > \mu_0(\alpha)$ (case B1). If $\mu_M \leq \mu_0(\alpha)$, types M and L are pooled after history M : $q_M(M) = q_L(M)$ (case B2).*

- **Case B3**. *If $\mu_M < \mu^{**}(\alpha)$, the optimal contract pools types M and L in the first period: $q_M = q_L$. In the second period, after history H, the contract is separating and efficient. After histories M and L, types M and L are pooled across both histories: $q_M(j) = q_L(j)$ and $q_j(M) = q_j(L)$ for $j = M, L$.*

Propositions 6 and 7 provide a full characterization of the optimal contract that can be used to gain new insights on how types are optimally screened in dynamic environments that are not apparent in the models discussed in Section 3.4.[30]

How does the possibility of repeated interactions affect the structure of the optimal contract? It is imperative to note that, in contrast to the static model, binding global constraints are no longer synonymous with pooling alone. In regions $B1$ and $B2$, even though the global incentive constraint binds, there is complete separation of types in period 1. Region $B2$ interestingly, like in $A2$, has a strategic separation in period 1 followed by history dependent pooling in period 2, which we term dynamic pooling. Region $B3$ has pooling in period 1 and in period 2 after histories θ_M and θ_L (but not θ_H). The contract captures a loss of history in region $B3$- it is as if we are in a

[30]Table 1, presented in the appendix, enlists closed form solutions of the optimal quantities for the entire parameter space.

two-type model in period 1 and after histories M and L in period 2. Further, note that the principal never chooses to pool types in period 1 without also having some pooling in period 2.

Another lesson is that pooling of types at $t = 1$ is always lower in the dynamic model than in a static optimal contract. It is easy to verify that in a static model types are pooled only if $\mu_M \leq \underline{\mu}_M = \frac{\mu_L(1-\mu_L)}{1+\mu_L}$, which in the example presented above is equal to 0.15- the horizontal line in Figure 3.7. This condition, however, is irrelevant in a dynamic model. We have pooling at $t = 1$ only if $\mu_M < \mu^{**}(\alpha)$, and $\mu^{**}(\alpha) < \underline{\mu}_M$ for all $\alpha \in (1/3, 1)$. The reason that the region for pooling at $t = 1$ is strictly smaller with two periods than with one is fairly intuitive: in a dynamic environment the principal has an added instrument in the form of continuation value to screen the agent's types. Thus, the burden of the efficiency rent trade-off can be pushed into the future, spreading distortions over time. Yet full separation in the static model does not imply full separation over time in the dynamic model as is evident in region $A2$, with $\mu_M > \underline{\mu}_M$, but pooling in period 2.

3.7 Ironing, implementability and optimality

The results of the previous sections make clear that in order to solve for an optimal contract, the principal cannot generally use the first-order approach and limit the analysis to local incentive compatibility constraints. Without the first-order approach, we have no systematic way of simplifying the constraint set. This may make the analysis extremely complicated even from a numerical point of view. What does the optimal contract look like in general environments with large T and N? What kind of advice can we give to a seller who needs to design an optimal contract? In this section we show that there is a class of contracts that is relatively easy to characterize, and that induces a minimal loss (if any) on the principal's payoff precisely when the first-order approach fails, that is, when the agent's types are highly persistent. This class consists of contracts that are monotonic in the sense of Definition 2. In static environments the envelope formula plus monotonicity are necessary and sufficient for a contract to be implementable: if we ignore the monotonicity constraint, then the contract *must* be ironed out to make it monotonic, otherwise implementability fails (see Myerson [1981]). In a dynamic environment monotonicity is not necessary: it follows that if we impose monotonicity in the seller's problem, we guarantee implementability even if we ignore the global constraints, but we may obtain a suboptimal contract. The main result of this section is that as types' persistence converge to one, the optimal monotonic contract converges in probability to the optimal contract, and so the loss from focusing on this class of contracts converges to zero.

Define $\mathcal{M}$ as the set of *monotonic contracts*:

$$\mathcal{M} = \left\{ \mathbf{q} \;\middle|\; \begin{array}{c} q(\theta_i|h^{t-1}) \geq q(\theta_{i+1}|h^{t-1}), i < N, \text{ and } q(\theta_i|h^{t-1}) \geq q(\theta_i|\widehat{h}^{t-1}), \\ i = 1, ..., N, \forall h^{t-1} \text{ and } h^{t-1} \succeq \widehat{h}^{t-1} \end{array} \right\} \tag{3.19}$$

where, as before, $h^t \succeq \widehat{h}^t$ if $h^t_j \geq \widehat{h}^t_j \; \forall j \leq t$. It follows immediately from Proposition 2 that the optimal monotonic contract can be characterized by solving the following program:

$$\max_{\mathbf{q} \in \mathcal{M}} \left\{ \mathbb{E}\left[S(\mathbf{q})\right] - \sum_{i=0}^{N} \mu_i U^*(\theta_i, h^0; \mathbf{q}) \right\} \tag{3.20}$$

where $U^*(\theta_i, h^0; \mathbf{q})$ is given by the envelope formula (3.5). Problem (3.20), moreover, is sufficiently tractable to allow a partial characterization of the properties of its solution.

Proposition 8. *In the optimal monotonic allocation,* $q\left(\theta_t \left| h^{t-1}\right.\right) \leq \theta_t$ for any θ_t and h^{t-1}. *Moreover, for any arbitrarily small* $\varepsilon_1, \varepsilon_2 > 0$ *we have* $\Pr\left(\left|q\left(\theta_t \left| h^{t-1}\right.\right) - \theta_t\right| > \varepsilon_1\right) \leq \varepsilon_2$ *for* t *and* T *sufficiently large.*

The first part of the proposition establishes that, analogous to the static model, the optimal monotonic contract is uniformly downward distorted. The second part states that the contract converges to an efficient contract in probability.

How good is the optimal monotonic contract as an approximation of the optimal contract? Let $\mathbf{q} = \{q(h^t)\}_{h^t \in H}$ be an allocation and let $\mathbf{q}^{**} = \{q^{**}(h^t)\}_{h^t \in H}$ be the optimal allocation. Let I be the identity matrix that describes the transition matrix when types are perfectly correlated. As types become perfectly persistent, we must have that the transition matrix converges to I, i.e. $\alpha \to I$. We say that $\mathbf{q}$ converges in probability to the optimal allocation as types become perfectly persistent if $\lim_{\alpha \to I} \Pr(|q(h^t) - q^{**}(h^t)| \geq \varepsilon) = 0$ for any $\varepsilon > 0$.

Proposition 9. *For all μ, δ, T, and transition matrices, the optimal monotonic contract converges in probability to a contract that maximizes the seller's profits as types become perfectly persistent.*

This result implies that for any δ and T, as types become increasingly persistent the profit associated with the optimal monotonic contract converges to the profit in the optimal contract. The table in Figure 3.8 illustrates the loss of profits associated with the optimal monotonic contract in an example with 3 periods, 3 types and the Markov matrix used in Section 6. The loss is expressed as a percentage of the profit in the optimal contract. As can be seen, the approximation is quite good for all cases, with a loss of profit that is never higher than 0.06%. It is interesting to note the inverse-U relationship between losses and the level of persistence. As persistence increases, losses increase, peak and then come down again. The reason is simple. At $\alpha = 1/3$, the model is akin to the i.i.d. shock framework, where we know that the optimal contract is monotonic. At the other extreme, $\alpha = 1$, the optimal contract constitutes repetition of the static optimum which too is monotonic. As we increase α, the distortions vary and the probability of non-constant histories decreases. Thus, the loss in using monotonic contracts increases with the non-monotonicities only to be suppressed in probability by the increasing weight of constant histories along which the optimal monotonic allocation converges to the optimal allocation.

When simultaneously types' persistence, the discount factor and the length of the contract are high, Proposition 9 may not be sufficient to guarantee that the optimal monotonic contract is a good approximation for the seller. Even a contract that converges to the efficient contract as $\alpha \to I$ may perform very poorly as $\delta \to 1$ and $T \to \infty$ as well. For example, the repetition of the optimal static contract converges in probability to an optimal contract (as shown in Example 2): for any given α (even arbitrarily close to I), however, the difference in profits between this contract and the optimal contract becomes arbitrarily large as $\delta \to 1$ and $T \to \infty$.[31] The problem is that the contract may not converge to the efficient contract fast enough in α. Therefore, in general, the order of limits may matter when we allow both the probability of persistence and the discount factor to converge to one.

The following result shows that for the optimal monotonic contract, profits converges to the optimal level independently of the order of limits. Define $\pi_m(\alpha, \delta, T)$ and $\pi^*(\alpha, \delta, T)$ to be the expected average discounted profits corresponding to the optimal monotonic contract and the optimal contract.[32] Moreover, let $\pi_m(\alpha, \delta) = \lim_{T\to\infty} \pi_m(\alpha, \delta, T)$ and $\pi^*(\alpha, \delta) = \lim_{T\to\infty} \pi^*(\alpha, \delta)$ be the limit expected average profits as $T \to \infty$.[33]

Proposition 10. *When $\alpha \to I$ and $\delta \to 1$, the profits of the optimal monotonic*

[31] As proven in Battaglini [2005], the contract becomes efficient and the seller appropriates all the surplus as $\delta \to 1$ when types are imperfectly persistent. With the repetition of the optimal static contract, however, per period surplus is below the efficient level and only a fraction is appropriated by the seller. The difference in discounted profits, therefore, becomes arbitrarily large as $\delta \to 1$ and $T \to \infty$ when types are imperfectly persistent.

[32] If $\Pi_m(\alpha, \delta, T)$ and $\Pi^*(\alpha, \delta, T)$ are the expected discounted profits correspondeing to the optimal monotonic contract, then $\pi_m(\alpha, \delta, T) = (1-\delta)\Pi_m(\alpha, \delta, T)$ and $\pi^*(\alpha, \delta, T) = (1-\delta)\Pi^*(\alpha, \delta, T)$.

[33] This limit exists without loss of generality since $\pi_m(\alpha, \delta, T)$ and $\pi^*(\alpha, \delta, T)$ are bounded for any α and δ.

δ = 0.95	α						
	0.38	0.48	0.58	0.68	0.78	0.88	0.98
$\mu_H = 0.5$ $\mu_M = 0.1$	**0.01** 11.00	**0.01** 9.87	**0.02** 8.49	**0.02** 6.87	**0.01** 4.98	**0.01** 2.86	**0.00** 0.51
$\mu_H = 0.5$ $\mu_M = 0.2$	**0.01** 10.70	**0.02** 9.62	**0.04** 8.32	**0.06** 6.77	**0.06** 4.96	**0.04** 2.87	**0.01** 0.51
$\mu_H = 0.5$ $\mu_M = 0.3$	**0.01** 10.01	**0.01** 9.87	**0.02** 8.51	**0.03** 6.91	**0.03** 5.06	**0.02** 3.93	**0.01** 0.52
$\mu_H = 0.3$ $\mu_M = 0.1$	**0.01** 10.75	**0.01** 9.73	**0.01** 8.45	**0.02** 6.91	**0.02** 5.08	**0.01** 2.95	**0.00** 0.53
$\mu_H = 0.3$ $\mu_M = 0.2$	**0.01** 10.61	**0.01** 9.61	**0.01** 8.37	**0.03** 6.87	**0.04** 5.08	**0.03** 3.98	**0.01** 0.54
$\mu_H = 0.3$ $\mu_M = 0.3$	**0.01** 10.41	**0.01** 9.42	**0.01** 8.20	**0.02** 6.72	**0.02** 4.97	**0.02** 2.92	**0.01** 0.53

Figure 3.8: Percentage loss of optimal objective (monopolist's profit) by using monotonic contracts (in bold) and repetition of the static optimum.

contract converges to the profits of the optimal contract independent of the order of limits: $\lim_{\delta\to1}\lim_{\alpha\to I}\pi_m(\alpha,\delta)=\lim_{\delta\to1}\lim_{\alpha\to I}\pi^*(\alpha,\delta)$ and $\lim_{\alpha\to I}\lim_{\delta\to1}\pi_m(\alpha,\delta)=\lim_{\alpha\to I}\lim_{\delta\to1}\pi^*(\alpha,\delta)$.

The table in Figure 3.8 illustrates this point comparing the loss of profits of the optimal monotonic contract with the loss of profits obtained with the repetition of the optimal static contract: the loss can be higher than 10% of the optimal profits, even in this simple example with only 3 periods. Naturally larger losses should be expected with longer horizons.

The results of this section may be useful in applied work. As mentioned in the introduction, many works in the applied literature postulate that the first-order approach works. The risk is that the contracts thus characterized are not incentive compatible. Further in the most natural environments, this risk can not be fully resolved by numerical methods. To the extent that it is not possible to check *all* the incentive compatibility constraints, studying optimal monotonic contracts may be a more robust option, since it guarantees implementability and it is equal to the true optimal contract with high probability when types are highly persistent.

3.8 Related literature

Our paper is related to four main literatures. First, we have the traditional literature studying dynamic principal-agent models when the agent's type follows a stochastic process and the allocation is chosen in every period. The first paper to use the first-order approach to study dynamic models and state an associated "envelope formula" is Baron and Besanko [1984].[34] Their paper states the formula in general terms and shows it to be sufficient in two benchmark cases: when types are constant over time, in which case the optimal dynamic contract corresponds to a repetition of the static optimum; and when types' realizations are independently distributed over time, in which case the optimal contract is efficient starting from period 2.[35] Extensions of this approach to environment with imperfect correlation of types are presented by Besanko [1985], Laffont and Tirole [1990] and Battaglini [2005]. Besanko [1985] extends the analysis to an infinite horizon with continuous types following a AR(1) process; Laffont and Tirole [1990] focus on a two periods environment with two types. Battaglini [2005] extends the two types model to an infinite horizon.[36] The main

[34]See Section 3 for a discussion of the first-order approach and envelope formula. See Stole [2001], Laffont and Martimort [2002], Milgrom (2004), and Bolton and Dewatripont [2005] for general discussions of the envelope formula in the static case.

[35]See also Roberts (1982) and Townsend (1982) for dynamic principal-agent models in which types are serially uncorrelated.

[36]Other important contributions in the dynamic contracting literature are Dewatripont [1989], Hart and Tirole [1988], Rey and Salanie (1990), Rustichini and Wolinsky [1995], Biehl [2001], Battaglini [2007], Williams [2010], Bergeman and Valimaki [2010], Strulovici [2011], Athey and Segal [2013], Boleslavsky and Said [2013], Maestri [2013]. These papers however focus on different aspects of the problem and limit the analysis to environment that are quite different from ours. Hart and Tirole (1988) assumes that supply can have two values, zero or one. Rustichini and Wolinsky (1995) assume consumers are not strategic and ignore that future prices depend on their current actions. Dewatripont (1989), Rey and Salanie (1990), Battaglini [2007], Maestri [2013] and Strulovici [2011] focus on renegotiation.

contributions of these papers is in showing that the first-order approach is sufficient in their respective environments. Laffont and Tirole [1996], and more recently, Pavan, Segal and Toikka [2013] and Eso and Szentes [2013] have derived "envelope formulas" for continuous types applicable to more complex environments. Both of the latter papers build on Eso and Szentes [2007], where the principal-agent problem is transformed in to a problem in which the shocks are i.i.d. through an appropriate change in utility. Contrary to the previous literature, these papers are not focused on finding specific environments in which these envelope formulas are sufficient for incentive compatibility, leaving open the question of the general applicability of the first-order approach.

The second literature to which our paper is related is the literature on sequential screening started by Courty and Li [2000]. This literature studies environments in which the agent receives information gradually over time, but the allocation is determined only in the last period. The models in this literature have 2 stages: in the beginning of period 1, the agent receives an informative signal and the contract is signed at the end of this period, but no allocation is made; in the second period the type is revealed to the agent and the allocation takes place. Courty and Li is one of the first papers to clearly discuss the limitations of the first-order approach in dynamic environments: one of their main achievements is to identify environments in which the first-order approach can be applied in the class of problems that they study. More recently, Courty and Li's work has been extended in many directions. Eso and Szentes [2007] consider the case in which the seller can choose to voluntarily disclose information in the first period. They show that the agent does not receive private rents for the disclosure of information. Li and Shi [2013] show that discriminatory disclosure of information can be optimal when the amount of additional private

Bergemann and Valimaki [2010] and Athey and Segal [2013] study implementation of efficient allocations extending the pivot mechanism to dynamic environments.

information that the buyer can learn depends on his type.[37] Krahmer and Strausz [2013] argue that in this class of models the benefit of sequential screening is due to the joint relaxation of incentive and participation constraints. To solve their model, the authors propose an original approach to deal with global constraints that works in their environment with N types. In all these papers the key question is whether the contract must depend on the interim informative signal, or if it can depend only on the type revealed in the last stage. In our model, because the allocation is chosen in all periods, information must be disclosed in all periods.

Third, our paper is related to a recent literature devoted to the study of approximately optimal mechanisms in environments in which fully optimal mechanisms are hard to characterize (see Madarasz and Prat [2012], Chassang [2013] for recent contributions and Hartline [2012] for a summary of the computer science approach). While parts of this literature deal with more general environments than ours, the approach we adopt in Section 3.7 takes full advantage of the dynamic structure of the framework we study; this allows us to obtain an approximately optimal contract that guarantees incentive compatibility for all types at all histories.

Finally, there is a large and growing literature using the first-order approach to solve dynamic contracts in complex environments using numerical methods. Understanding the conditions for the applicability of the first-order approach with discrete types seems particularly important in these exercises. Even when using models with continuous types, these papers typically compute the equilibrium policies and verify incentive compatibility using discretized approximations.[38] When discrete approximations are not used to construct the first-order optimal contract, incentive compatibility

[37] For other recent models of infomation disclosure see, amongst others, Hoffman and Inderst [2011], Inderst and Ottaviani [2012].

[38] This is the case, for example, in Kapicka [2013], Farhi and Werning [2013], and Golosov et al. [2013] who study models of intertemporal consumption smoothing using numerical methods.

is verified numerically on a grid of points.[39] The envelope formula presented in our paper provides an exact formula for discrete types that can be used to compute the first-order optimal contract and to verify incentive compatibility directly without approximations.

3.9 Conclusion

In this paper we have studied a simple principal-agent model in which the agent's type is private information and follows a Markov process. We have presented four sets of results. First, following the standard approach in the literature, we have studied the optimal contract when only local incentive constraints are considered. We have shown that the agent's equilibrium rents can be represented purely as a function of the allocation through a dynamic version of the so called "envelope formula." Moreover, as in the static model, the envelope formula and a natural monotonicity condition on the allocation guarantee that the contract is implementable. Although this condition is only sufficient and quite strong, it is verified for virtually all the natural environments in which the optimal dynamic contract has been characterized in the existing literature.

Second, and most importantly, we have shown that the environments for which the envelope formula is sufficient to characterize the optimal dynamic contract are quite special. In general, even in the simplest examples, the allocation is not monotonic. Thus, for high persistence and sufficiently long time horizons global incentive constraints generically bind. Moreover, numerical examples show that

[39]Exceptions are Zhang [2009] and Williams [2011] who use continuous time methods to avoid discrete approximation of the policy functions. Zhang [2009] and Williams [2011] verify that the conditions for the first order approach are satisfied in their model. Zhang [2009] however, limits the analysis to a two types model; and Williams [2011] limits the set of possible deviations available to the agent (who can report only incomes lower or equal to the true income).

moderate levels of persistence are sufficient to violate the first-order approach.

Third, to gain insight on how the optimal contract looks like when the first-order approach doesn't work, we have characterized it in a simple case with three types and two periods. We show that the optimal contract is characterized by *dynamic pooling*: strategic, state contingent treatment of types in which types may be initially separated, but then be pooled conditioned on particular histories.

Finally, we have shown that some insights in general environments with many types and periods can be gained by studying a simple class of suboptimal contracts: monotonic contracts, in which non-monotonicities in the allocation are "ironed" out. The appeal of optimal monotonic contracts is derived from the fact that it converges in probability to the optimal contract as the persistence of types converges to one, that is precisely when the first-order approach tends to fail.

The analysis suggests a number of important research questions. The characterization of the optimal contract with three types and two periods suggests that state dependent pooling of types plays an important role in dynamic screening. The example suggests a number of features that one naturally expects to hold in more general environments as well. The analysis in Section 7, moreover, suggests that even when it is not possible to fully characterize the optimal contract, useful insights can be gained by studying contracts that are approximately optimal. We leave the further development of these ideas for future research.

3.10 Appendix

3.10.1 Proof of Lemma 1 and Corollary 1

We first show that all the constraints in the relaxed problem can be assumed to hold as equalities.

Lemma A1. *In a FO-relaxed problem:* $IR_N(h^{t-1})$ *can be assumed to hold as equality* *for all* $h^{t-1} \in H^{t-1}$*;* $IC_{i,i+1}(h^{t-1})$ *can be assumed to hold as an equality for all* $h^{t-1} \in H^{t-1}$ *and* $i = 0, 1, ..., N-1$.

Proof. We proceed in two steps:

Step 1. Suppose that $U(\theta_N|h^{t-1}) = \epsilon > 0$ for some h^{t-1}. If $t = 1$, then decreasing $U(\theta_i|h^0)$ by ϵ for all i does not violate any constraints and increases the monopolist's profit. If $t > 1$, fix h^{t-1} and decrease $U(\theta_i|h^{t-1})$ by ϵ for all θ_i. This does not change any of the constraints and keeps the profit of the monopolist the same.

Step 2. Suppose that $IC_{i,i+1}(h^{t-1})$ does not hold as an equality for some $h^{t-1} \in H^{t-1}$ and $i = 0, 1, ..., N-1$. Then, decrease $U(\theta_k|h^{t-1})$ by ϵ for each $k \leq i$. If $t = 1$, all the constraints are still satisfied and the monopolist's profit is strictly higher, giving a contradiction. If $t > 1$, this change does not affect any constraint except $IC_{j-1,j}(h^{t-2})$, where θ_j is such that $h^{t-1} = (h^{t-2}, \theta_j)$. The right hand side of $IC_{j-1,j}(h^{t-2})$ is *reduced* by $\delta \sum_{k \leq i}(\alpha_{(j-1)k} - \alpha_{jk})\epsilon = \delta\Delta F\left(\theta_{i+1}|\theta_j\right)\epsilon \geq 0$, where the last inequality follows from first order stochastic dominance. Now, repeat the same procedure, decreasing $U(\theta_k|h^{t-2})$ by $\delta\Delta F\left(\theta_{i+1}|\theta_j\right)\epsilon$ for each $k \leq j-1$. We can keep reducing utility vectors backward till the first period, unless h^{t-1} contains θ_0, in which case the backward iteration ends there, to deduce a strictly greater increase in the monopolist's profit. Thus, the changes do not violate any of the constraints and keep the profit of the monopolist larger than or equal to before the change. ■

We can now prove Lemma 1 and Corollary 1 together. We shall proceed by (backward) induction on t. Note that at $t = T$, Lemma A1 implies:

$$U(\theta_N|h^{T-1}) = 0 \text{ and } U(\theta_i|h^{T-1}) = \sum_{l=1}^{N-i} \Delta u(\theta_{i+l}|h^{T-1}; \mathbf{q}) \; \forall i \leq N-1. \tag{3.21}$$

where $\Delta u(\theta_{i+1}|h^{t-1}; \mathbf{q})$ is defined by

$$\Delta u(\theta_{i+1}|h^{t-1}; \mathbf{q}) = u(\theta_i, q(\theta_{i+1}|h^{t-1})) - u(\theta_{i+1}, q(\theta_{i+1}|h^{t-1})).$$

Similarly, for $t = T-1$, we have for $i \leq N-1$:

$$\begin{aligned}
U(\theta_i|h^{T-2}) &= \Delta u(\theta_{i+1}|h^{T-2};\mathbf{q}) + U(\theta_{i+1}|h^{T-2}) + \delta\sum_{k=0}^{N}(\alpha_{ik} - \alpha_{(i+1)k})U(\theta_k|h^{T-2},\theta_{i+1}) \\
&= \sum_{n=1}^{N-i}\Big[\Delta u(\theta_{i+n}|h^{T-2};\mathbf{q}) + \delta\sum_{k=0}^{N}(\alpha_{(i+n-1)k} - \alpha_{(i+n)k})U(\theta_k|h^{T-2},\theta_{i+n})\Big] \\
&= \sum_{n=1}^{N-i}\Big[\Delta u(\theta_{i+n}|h^{T-2};\mathbf{q}) + \\
&\qquad \delta\sum_{k=0}^{N-1}(\alpha_{(i+n-1)k} - \alpha_{(i+n)k})\sum_{l=1}^{N-k}\Delta u(\theta_{k+l}|h^{T-2},\theta_{i+n};\mathbf{q}))\Big]
\end{aligned}$$

Now, let

$$\begin{aligned}
&\sum_{k=0}^{N-1}(\alpha_{(i+n-1)k} - \alpha_{(i+n)k})\sum_{l=1}^{N-k}\Delta u(\theta_{k+l}|h^{T-2},\theta_{i+n};\mathbf{q})) \\
&\qquad = \sum_{k=0}^{N-1}\left(\alpha_{(i+n-1)k} - \alpha_{(i+n)k}\right)\sum_{l=1}^{N-k}Q_{k+l}, \qquad (3.22)
\end{aligned}$$

where $Q_j = \Delta u(\theta_j|h^{T-2},\theta_{i+n};\mathbf{q}))$ for for any type θ_j. The right hand side of (3.22) can be written as:

$$\sum_{k=0}^{N-1}\left(\alpha_{(i+n-1)k} - \alpha_{(i+n)k}\right)\sum_{l=1}^{N-k}Q_{k+l} = \left[\begin{array}{c} \left(\alpha_{(i+n-1)0} - \alpha_{(i+n)0}\right)(Q_1 + ... + Q_N) \\ + \left(\alpha_{(i+n-1)1} - \alpha_{(i+n)1}\right)(Q_2 + ... + Q_N) \\ +... + \left(\alpha_{(i+n-1)(N-1)} - \alpha_{(i+n)(N-1)}\right)Q_N \end{array}\right]$$

Rearranging the terms, we have:

$$
\begin{aligned}
&\sum_{k=0}^{N-1}\left(\alpha_{(i+n-1)k}-\alpha_{(i+n)k}\right)\sum_{l=1}^{N-k}Q_{k+l} \\
= &\left[\begin{array}{c}
\left(\alpha_{(i+n-1)0}-\alpha_{(i+n)0}\right)Q_1 \\
+\left(\left(\alpha_{(i+n-1)0}+\alpha_{(i+n-1)1}\right)-\left(\alpha_{(i+n)0}+\alpha_{(i+n)1}\right)\right)Q_2 \\
+...+\left(\left(\alpha_{(i+n-1)0}+...+\alpha_{(i+n-1)(N-1)}\right)-\left(\alpha_{(i+n)0}+...+\alpha_{(i+n)(N-1)}\right)\right)Q_N
\end{array}\right] \\
= &\sum_{k=1}^{N}\Delta F(\theta_k|\theta_{i+n})Q_k
\end{aligned}
$$

where, we recall, $\Delta F\left(\theta_j\left|\theta_i\right.\right)=F\left(\theta_j\left|\theta_i\right.\right)-F\left(\theta_j\left|\theta_{i-1}\right.\right)$. This implies that:

$$
\begin{aligned}
&\sum_{k=0}^{N-1}(\alpha_{(i+n-1)k}-\alpha_{(i+n)k})\sum_{l=1}^{N-k}\Delta u(\theta_{k+l}|h^{T-2},\theta_{i+n};\mathbf{q})) \\
&\quad=\sum_{k=1}^{N}\Delta F(\theta_k\left|\theta_{i+n}\right.)\Delta u(\theta_k|h^{T-2},\theta_{i+n};\mathbf{q}))
\end{aligned}
$$

It follows that we can write:

$$
\begin{aligned}
U(\theta_i|h^{T-2}) &= \sum_{n=1}^{N-i}\left[\Delta u(\theta_{i+n}|h^{T-2};\mathbf{q})+\delta\sum_{k=1}^{N}\Delta F(\theta_k\left|\theta_{i+n}\right.)\Delta u(\theta_k|h^{T-2},\theta_{i+n};\mathbf{q}))\right] \\
&= \sum_{n=1}^{N-i}\left[\Delta u(\theta_{i+n}|h^{T-2};\mathbf{q})\right. \\
&\quad\left.+\sum_{\hat{h}\in\widehat{H}(h^{T-2},\theta_{i+n})}\sum_{\tau>T-1}\delta^{\tau-T-1}\prod_{k=T}^{\tau}\Delta F\left(\hat{h}_k\left|\hat{h}_{k-1}\right.\right)\Delta u(\hat{h}_\tau|\hat{h}^{\tau-1};\mathbf{q}))\right]
\end{aligned}
\tag{3.23}
$$

where, we reacall, $\widehat{H}(h^t)$ is the set of histories following h^t in which all realizations after t are lower than θ_0.

It is easy to see that (3.21) and (3.23) prove the statement in Corollary 1 and in Lemma 1 respectively for $t=T$ and $t=T-1$. We therefore conclude that our hypothesis holds for $t\geq T-1$. Next, suppose it holds for $t+1$ where $t\geq T-2$. We want to show that it holds for t. We have,

$$U(\theta_i|h^{t-1}) = \Delta u\left(\theta_{i+1}|h^{t-1};\mathbf{q}\right) + U\left(\theta_{i+1}|h^{t-1}\right) + \delta\sum_{k=0}^{N}(\alpha_{ik} - \alpha_{(i+1)k})U(\theta_k|h^{t-1},\theta_{i+1}) \tag{3.24}$$

$$= \sum_{n=1}^{N-i}\left[\Delta u\left(\theta_{i+n}|h^{t-1};\mathbf{q}\right) + \delta\sum_{k=0}^{N}(\alpha_{(i+n-1)k} - \alpha_{(i+n)k})U\left(\theta_k|h^{t-1},\theta_{i+n}\right)\right]$$

$$= \sum_{n=1}^{N-i}\left[\Delta u\left(\theta_{i+n}|h^{t-1};\mathbf{q}\right) + \delta\sum_{k=0}^{N-1}(\alpha_{(i+n-1)k} - \alpha_{(i+n)k})\right.$$

$$\sum_{m=1}^{N-k}\left(\Delta u\left(\theta_{k+m}|h^{t-1},\theta_{i+n};\mathbf{q}\right) + \right.$$

$$\left.\left.\sum_{\hat{h}\in\widehat{H}(h^{t-1},\theta_{i+n},\theta_{k+m})}\sum_{\tau>t+1}\delta^{\tau-(t+1)}\prod_{\iota=t+2}^{\tau}\Delta F\left(\hat{h}_\iota\left|\hat{h}_{\iota-1}\right.\right)\Delta u\left(\hat{h}_\tau|\hat{h}^{\tau-1};\mathbf{q}\right)\right)\right],$$

where the third equality follows from the induction hypothesis. Now,

$$\sum_{k=0}^{N-1}(\alpha_{(i+n-1)k} - \alpha_{(i+n)k})\sum_{m=1}^{N-k}\Delta u(\theta_{k+m}|h^{t-1},\theta_{i+n};\mathbf{q})$$

$$= \sum_{k=1}^{N}\Delta F(\theta_k\,|\theta_{i+n})\Delta u(\theta_k|h^{t-1},\theta_{i+n};\mathbf{q}), \tag{3.25}$$

and,

$$\delta\sum_{k=0}^{N-1}(\alpha_{(i+n-1)k} - \alpha_{(i+n)k})\sum_{m=1}^{N-k}\sum_{\hat{h}\in\widehat{H}(h^{t-1},\theta_{i+n},\theta_{k+m})}\sum_{\tau>t+1}$$

$$\delta^{\tau-(t+1)}\prod_{\iota=t+2}^{\tau}\Delta F\left(\hat{h}_\iota\left|\hat{h}_{\iota-1}\right.\right)\Delta u\left(\hat{h}_\tau|\hat{h}^{\tau-1};\mathbf{q}\right) = \delta\sum_{k=0}^{N-1}(\alpha_{(i+n-1)k} - \alpha_{(i+n)k})\sum_{m=1}^{N-k}Q_{k+m},$$

where, $Q_l = \sum_{\hat{h}\in\widehat{H}(h^{t-1},\theta_{i+n},\theta_l)}\sum_{\tau>t+1}\delta^{\tau-(t+1)}\prod_{\iota=t+2}^{\tau}\Delta F\left(\hat{h}_\iota\left|\hat{h}_{\iota-1}\right.\right)\Delta u\left(\hat{h}_\tau|\hat{h}^{\tau-1};\mathbf{q}\right)$. As before, after some algebraic manipulation, this becomes:

$$\delta\sum_{k=0}^{N-1}(\alpha_{(i+n-1)k} - \alpha_{(i+n)k})\sum_{m=1}^{N-k}Q_{k+m} = \delta\sum_{k=1}^{N}\Delta F\left(\theta_k|\theta_{i+n}\right)Q_k \tag{3.26}$$

Combining (3.25) and (3.26) we obtain:

$$\delta \sum_{k=1}^{N} \Delta F\left(\theta_k|\theta_{i+n}\right)\left[\Delta u(\theta_k|h^{t-1},\theta_{i+n};\mathbf{q}) + Q_k\right] \tag{3.27}$$
$$= \delta \sum_{k=1}^{N} \Delta F\left(\theta_k|\theta_{i+n}\right)\Bigg[\Delta u(\theta_k|h^{t-1},\theta_{i+n};\mathbf{q}) +$$
$$\sum_{\hat{h}\in\widehat{H}(h^{t-1},\theta_{i+n},\theta_k)} \sum_{\tau>t+1} \delta^{\tau-(t+1)} \prod_{\iota=t+2}^{\tau} \Delta F\left(\hat{h}_\iota \middle| \hat{h}_{\iota-1}\right) \Delta u\left(\hat{h}_\tau|\hat{h}^{\tau-1};\mathbf{q}\right)\Bigg]$$
$$= \sum_{\hat{h}\in\widehat{H}(h^{t-1},\theta_{i+n})} \sum_{\tau>t} \delta^{\tau-t} \prod_{\iota=t+1}^{\tau} \Delta F\left(\hat{h}_\iota \middle| \hat{h}_{\iota-1}\right) \Delta u\left(\hat{h}_\tau|\hat{h}^{\tau-1};\mathbf{q}\right)$$

Combining (3.24) and (3.27), we obtain:

$$U\left(\theta_i|h^{t-1}\right) = \sum_{n=1}^{N-i}\Bigg[\Delta u\left(\theta_{i+n}|h^{t-1};\mathbf{q}\right) +$$
$$\sum_{\hat{h}\in\widehat{H}(h^{t-1},\theta_{i+n})} \sum_{\tau>t} \delta^{\tau-t} \prod_{\iota=t+1}^{\tau} \Delta F\left(\hat{h}_\iota \middle| \hat{h}_{\iota-1}\right) \Delta u(\hat{h}_\tau|\hat{h}^{\tau-1};\mathbf{q})\Bigg].$$

Note that:

$$\Delta u(\theta_{i+1}|h^{t-1};\mathbf{q}) = u(\theta_i, q(\theta_{i+1}|h^{t-1})) - u(\theta_{i+1}, q(\theta_{i+1}|h^{t-1})) = \int_{\theta_{i+1}}^{\theta_i} u_\theta(x, q(\theta_{i+1}|h^{t-1}))dx$$

It follows that we have:

$$U\left(\theta_i|h^{t-1}\right) = \sum_{n=1}^{N-i}\Bigg[\int_{\theta_{i+n}}^{\theta_{i+n-1}} u_\theta(x, q(\theta_{i+n}|h^{t-1}))dx +$$
$$\sum_{\hat{h}\in\widehat{H}(h^{t-1},\theta_{i+n})} \sum_{\tau>t} \delta^{\tau-t} \prod_{\iota=t+1}^{\tau} \Delta F\left(\hat{h}_\iota \middle| \hat{h}_{\iota-1}\right) \int_{\hat{h}_\tau}^{\hat{h}_\tau+\Delta\theta} u_\theta(x, q(\hat{h}_\tau|\hat{h}^{\tau-1}))dx\Bigg].$$

This proves Corollary 1. Subtracting $U(\theta_{i+1}|h^{t-1})$ and dividing by $\Delta\theta$ from the above expression gives us Lemma 1.

3.10.2 Proof of Proposition 2

Recall that $\Delta U(\theta_k\,|h^{t-1},\theta_i) = U(\theta_k\,|h^{t-1},\theta_i) - U(\theta_k\,|h^{t-1},\theta_{i+1})$. We start with some useful lemmas.

Lemma A2. *If $q(\theta_i|h^{t-1})$ and $\Delta U(\theta_k\,|h^{t-1})$ are non increasing in, respectively, i and k for any h^{t-1}, then (3.5) implies that local upward incentive compatibility constraints are satisfied.*

Proof. Condition (3.5) implies that local downward constraints, $IC_{i,i+1}(h^{t-1})$, hold as equality for any i and h^{t-1}, that is:

$$U(\theta_i|h^{t-1}) = U(\theta_{i+1}|h^{t-1}) + \Delta\theta q(\theta_{i+1}|h^{t-1}) + \delta\sum_{k=0}^{N}\left(\alpha_{ik} - \alpha_{(i+1)k}\right)U(\theta_k|h^{t-1},\theta_{i+1}).$$

Thus,

$$\begin{aligned}
U(\theta_{i+1}|h^{t-1}) - U(\theta_i|h^{t-1}) &= -\Delta\theta q(\theta_{i+1}|h^{t-1}) - \delta\sum_{k=0}^{N}\left(\alpha_{ik} - \alpha_{(i+1)k}\right)U(\theta_k|h^{t-1},\theta_{i+1})) \\
&= -\Delta\theta q(\theta_i|h^{t-1}) + \delta\sum_{k=0}^{N}\left(\alpha_{(i+1)k} - \alpha_{ik}\right)U(\theta_k|h^{t-1},\theta_i) \\
&+ \Delta\theta\left(q(\theta_i|h^{t-1}) - q(\theta_{i+1}|h^{t-1})\right) + \\
&\quad \delta\sum_{k=0}^{N}\left(\alpha_{ik} - \alpha_{(i+1)k}\right)\Delta U(\theta_k|h^{t-1},\theta_i) \\
&\geq -\Delta\theta q(\theta_i|h^{t-1}) + \delta\sum_{k=0}^{N}\left(\alpha_{(i+1)k} - \alpha_{ik}\right)U(\theta_k|h^{t-1},\theta_i),
\end{aligned}$$

where the last inequality follows from the fact that $q(\theta_i|h^{t-1})$ is non increasing in i and $\sum_{k=0}^{N}(\alpha_{ik} - \alpha_{(i+1)k})\Delta U(\theta_k\,|h^{t-1},\theta_i) \geq 0$. The second observation follows from the fact that $\Delta U(\theta_k\,|h^{t-1},\theta_i)$ is non increasing in k, and that α_{ik} first order stochastically dominates $\alpha_{(i+1)k}$. Thus, $IC_{i+1,i}(h^{t-1})$ holds. ■

Lemma A3. *If $q(\theta_i|h^{t-1})$ and $\Delta U(\theta_k\,|h^{t-1})$ are non increasing in, respectively, i and k for any h^{t-1} and (3.5) holds, then the local incentive compatibility constraints* *imply the global incentive compatibility constraints.*

Proof. We show that $IC_{i,i+2}(h^{t-1})$ holds. The envelope formula (3.5) is equivalent to assuming that all the local downward incentive compatibility constraints are satisfied as equalities. From $IC_{i,i+1}(h^{t-1})$ and $IC_{i+1,i+2}(h^{t-1})$ we have:

$$\begin{aligned}
&U(\theta_i|h^{t-1}) - U(\theta_{i+2}|h^{t-1}) \\
&\quad= \left[U(\theta_i|h^{t-1}) - U(\theta_{i+1}|h^{t-1})\right] + \left[U(\theta_{i+1}|h^{t-1}) - U(\theta_{i+2}|h^{t-1})\right] \\
&\quad= \Delta\theta q(\theta_{i+1}|h^{t-1}) + \delta \sum_{k=0}^{N} \left(\alpha_{ik} - \alpha_{(i+1)k}\right) U(\theta_k|h^{t-1}, \theta_{i+1}) \\
&\quad+ \Delta\theta q(\theta_{i+2}|h^{t-1}) + \delta \sum_{k=0}^{N} \left(\alpha_{(i+1)k} - \alpha_{(i+2)k}\right) U(\theta_k|h^{t-1}, \theta_{i+2}).
\end{aligned}$$

It follows that:

$$\begin{aligned}
& U(\theta_i|h^{t-1}) - U(\theta_{i+2}|h^{t-1}) \\
= \;& 2\Delta\theta q(\theta_{i+2}|h^{t-1}) + \delta \sum_{k=0}^{N} \left(\alpha_{ik} - \alpha_{(i+2)k}\right) U(\theta_k|h^{t-1}, \theta_{i+2}) \\
& +\Delta\theta \left(q(\theta_{i+1}|h^{t-1}) - q(\theta_{i+2}|h^{t-1})\right) + \delta \sum_{k=0}^{N} \left(\alpha_{ik} - \alpha_{(i+1)k}\right) \Delta U(\theta_k|h^{t-1}, \theta_{i+1}) \\
\geq \;& \Delta\theta q(\theta_{i+2}|h^{t-1}) + \delta \sum_{k=0}^{N} \left(\alpha_{(i+1)k} - \alpha_{(i+2)k}\right) U(\theta_k|h^{t-1}, \theta_{i+2}),
\end{aligned}$$

where the last inequality follows from the fact that $q(\theta_i|h^{t-1})$ is non increasing in i and $\sum_{k=0}^{N}(\alpha_{ik} - \alpha_{(i+1)k})\Delta U(\theta_k\,|h^{t-1}, \theta_i) \geq 0$. As in the previous lemma, the second observation follows from the fact that $\Delta U(\theta_k\,|h^{t-1}, \theta_i)$ is non increasing in k, and that α_{i+1k} first order stochastically dominates $\alpha_{(i+2)k}$. Thus, $IC_{i,i+2}(h^{t-1})$ holds. Similarly we can show that $IC_{i,i+l}(h^{t-1})$ holds for all $l \leq N - i$. Therefore, all global downward incentive constraints are satisfied. In an analogous fashion, we can show that all upward global incentive constraints are satisfied. ■

Given the lemmas presented above, Proposition 2 is proven if we establish that when the allocation is monotonic as defined in Definition 2, then $q(\theta_i|h^{t-1})$ and $\Delta U(\theta_k\,|h^{t-1},\theta_i)$ are non increasing in i for any h^{t-1}. The fact that $q(\theta_i|h^{t-1})$ is non increasing in i for any h^{t-1} is an immediate consequence of the monotonicity. The fact that $\Delta U(\theta_k\,|h^{t-1},\theta_i)$ is non increasing in i for any h^{t-1} is established by the following result.

Lemma A4. *If the allocation is monotonic as in Definition 2, then* $\Delta U(\theta_k\,|h^{t-1})$ *is non increasing in* k $\forall h^{t-1}$.

Proof. Note first that $U(\theta_N\,|h^{t-1},\theta_i) = U(\theta_N\,|h^{t-1},\theta_{i+1}) = 0$, so $\Delta U(\theta_N\,|h^{t-1},\theta_i) = 0$. By Lemma 1, we have:

$$U(\theta_{N-1}|h^{t-1},\theta_i) = \int_{\theta_N}^{\theta_{N-1}} u_\theta(x, q(\theta_N|h^{t-1},\theta_i))dx + \sum_{\hat{h}\in\widehat{H}(h^{t-1},\theta_i,\theta_{N-1})}\sum_{\tau>t+1}\delta^{\tau-t-1}\left[\begin{array}{c}\prod_{\iota=t+2}^{\tau}\Delta F\left(\hat{h}_\iota\,\middle|\hat{h}_{\iota-1}\right)\\ \cdot\int_{\hat{h}_\tau}^{\hat{h}_\tau+\Delta\theta} u_\theta(x,q(\hat{h}_\tau|\hat{h}^{\tau-1}))dx\end{array}\right] \tag{3.28}$$

It is useful to write this expression with a different notation. Let $\widehat{H}_t(i)$ be set of realizations of length $T-t$ that start with the first element equal to θ_i (we denote ${}_th$ is the typical element of $\widehat{H}_t(i)$, so ${}_th_1 = \theta_i$). A history $h^\tau \in \widehat{H}\,(h^t)$ with $(t+1)$-th element equal to θ_i ($h^\tau_{t+1} = \theta_i$) is then $h^\tau = \{h^t,{}_t\,h^{\tau-t}\}$ for ${}_th \in \widehat{H}_t(i)$ (by convention we write $h^t = \{h^t,{}_t\,h^0\}$). We can then write:

$$U(\theta_{N-1}\left|h^{t-1},\theta_i\right.) = \int_{\theta_N}^{\theta_{N-1}} u_\theta(x, q(\theta_N|h^{t-1},\theta_i))dx +$$

$$\sum_{{}_th\in\widehat{H}_t(N-1)}\sum_{\tau>t+1}\delta^{\tau-t-1}\left[\begin{array}{c}\prod_{l=t+2}^{\tau}\Delta F\left({}_th_l\,|{}_th_{l-1}\right)\cdot\\ \int_{{}_th_\tau}^{{}_th_\tau+\Delta\theta} u_\theta(x,q({}_th_\tau|h^{t-1},\theta_{i,t}\,h^{\tau-t-1}))\end{array}\right] \tag{3.29}$$

Similarly we can write:

$$U(\theta_{N-1}\left|h^{t-1},\theta_{i+1}\right.)\int_{\theta_N}^{\theta_{N-1}} u_\theta(x,q(\theta_N|h^{t-1},\theta_{i+1}))dx+$$

$$\sum_{{}_t h\in\widehat{H}_t(N-1)}\sum_{\tau>t+1}\delta^{\tau-t-1}\left[\begin{array}{c}\prod_{l=t+2}^{\tau}\Delta F\left({}_t h_l\left|{}_t h_{l-1}\right.\right)\cdot \\ \int_{{}_t h_\tau}^{{}_t h_\tau+\Delta\theta} u_\theta(x,q({}_t h_\tau|h^{t-1},\theta_{i+1},{}_t h^{\tau-t-1}))dx\end{array}\right] \tag{3.30}$$

Therefore we have:

$$\begin{aligned}&\Delta U(\theta_{N-1}\left|h^{t-1},\theta_i\right.)= \\ &\int_{\theta_N}^{\theta_{N-1}}\left[u_\theta(x,q(\theta_N|h^{t-1},\theta_i))-u_\theta(x,q(\theta_N|h^{t-1},\theta_{i+1}))\right]dx \\ &+\sum_{{}_t h\in H_t(N-1)}\sum_{\tau>t+1}\delta^{\tau-t-1}\left[\begin{array}{c}\prod_{l=t+2}^{\tau}\Delta F\left({}_t h_l\left|{}_t h_{l-1}\right.\right) \\ \cdot\int_{{}_t h_\tau}^{{}_t h_\tau-\Delta\theta}\left[\begin{array}{c}u_\theta(x,q({}_t h_\tau|h^{t-1},\theta_{i},{}_t h^{\tau-t-1})) \\ -u_\theta(x,q({}_t h_\tau|h^{t-1},\theta_{i+1},{}_t h^{\tau-t-1}))\end{array}\right]dx\end{array}\right]\end{aligned}$$

Note that by monotonicity, we must have $q(\theta_N|h^{t-1},\theta_i)-q(\theta_N|h^{t-1},\theta_{i+1})\geq 0$ and

$$q({}_t h_\tau|h^{t-1},\theta_{i},{}_t h^{\tau-t-1})-q({}_t h_\tau|h^{t-1},\theta_{i+1},{}_t h^{\tau-t-1})\geq 0$$

The above condition plus the single crossing condition (Assumption 1) imply that $\Delta U(\theta_{N-1}\left|h^{t-1},\theta_i\right.)\geq\Delta U(\theta_N\left|h^{t-1},\theta_i\right.)$. Assume now that $\Delta U(\theta_j\left|h^{t-1},\theta_i\right.)$ is monotonic in j for $j\geq m$. We show below that $\Delta U(\theta_{m-1}\left|h^{t-1},\theta_i\right.)\geq\Delta U(\theta_m\left|h^{t-1},\theta_i\right.)$, the result then follows from induction. Applying Lemma 1 and using the notation developed above, we have:

$$
\begin{aligned}
& \Delta U(\theta_{m-1} \left| h^{t-1}, \theta_i \right.) \\
= \; & \Delta U(\theta_m \left| h^{t-1}, \theta_i \right.) + \int_{\theta_N}^{\theta_{N-1}} \left[u_\theta(x, q(\theta_m | h^{t-1}, \theta_i)) - u_\theta(x, q(\theta_m | h^{t-1}, \theta_{i+1})) \right] dx \\
& + \sum_{{}_th \in H_t(m-1)} \sum_{\tau > t+1} \delta^{\tau-t-1} \left[\begin{array}{c} \prod_{l=t+2}^{\tau} \Delta F\left({}_th_l \left| {}_th_{l-1} \right. \right) \\ \cdot \int_{{}_th_\tau}^{{}_th_\tau - \Delta\theta} \left[\begin{array}{c} u_\theta(x, q({}_th_\tau | h^{t-1}, \theta_{i,t}\, h^{\tau-t-1})) \\ -u_\theta(x, q({}_th_\tau | h^{t-1}, \theta_{i+1,t}\, h^{\tau-t-1})) \end{array} \right] dx \end{array} \right]
\end{aligned}
$$

Thus, the single crossing condition and monotonicity of the allocation imply

$$\Delta U(\theta_{m-1} \left| h^{t-1}, \theta_i \right.) \geq \Delta U(\theta_m \left| h^{t-1}, \theta_i \right.).$$

■

3.10.3 Proofs of Propositions 3

For $0 < i < N$, define $\Psi_i(f_\alpha)$ as:

$$
\begin{aligned}
\Psi_i(f_\alpha) &= \left[\frac{\Delta F_\alpha(\theta_i | \theta_{i+1})}{f_\alpha(\theta_i | \theta_{i+1})} \cdot \frac{\Delta F_\alpha(\theta_{i+1} | \theta_i)}{f_\alpha(\theta_{i+1} | \theta_i)} \right] \\
&= \frac{\sum_{k=i+1}^{N} \left[f_\alpha(\theta_k | \theta_i) - f_\alpha(\theta_k | \theta_{i-1}) \right]}{f_\alpha(\theta_{i+1} | \theta_i)} \cdot \frac{\sum_{k=i}^{N} \left[f_\alpha(\theta_k | \theta_{i+1}) - f_\alpha(\theta_k | \theta_i) \right]}{f_\alpha(\theta_i | \theta_{i+1})}
\end{aligned}
$$

In Lemma A5 we prove that if:

$$\lim_{\alpha \to 1} \Psi_i(f_\alpha) \neq 1 \text{ for some } i \in (0, N) \tag{3.31}$$

then the optimal contract is not monotonic as $\alpha \to 1$. In Lemma A6 we prove that condition (3.31) is generically satisfied.

Lemma A5. *For any μ, δ, $|\Theta| > 2$, $T > 2$, if $\lim_{\alpha\to 1}\Psi_i(f_\alpha) \neq 1$ for some $i \in (0, N)$, then there is an $\alpha^* < 1$ such that the FO-optimal contract is not monotonic for any $\alpha > \alpha^*$.*

Proof. Suppose first $D = \lim_{\alpha\to 1}\Psi_i(f_\alpha) < 1$. We show that $\lim_{\alpha\to 1} q(\theta_i|\theta_i, \theta_{i+1}) > \lim_{\alpha\to 1} q(\theta_i|\theta_i, \theta_i)$. Note that

$$s_q(\theta_i, q(\theta|\theta_i, \theta_i)) \leq \frac{1-\sum_{k=i}^N \mu_k}{\mu_i} \cdot \left[\frac{\Delta F_\alpha(\theta_i|\theta_i)}{f_\alpha(\theta_i|\theta_i)} \cdot \frac{\Delta F_\alpha(\theta_i|\theta_i)}{f_\alpha(\theta_i|\theta_i)}\right] \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q(\theta_i|\theta_i, \theta_i)dx$$

and

$$s_q(\theta_i, q(\theta_i|\theta_i, \theta_{i+1})) \leq$$
$$\frac{1-\sum_{k=i}^N \mu_k}{\mu_i} \cdot \left[\frac{\Delta F_\alpha(\theta_i|\theta_{i+1})}{f_\alpha(\theta_i|\theta_{i+1})} \cdot \frac{\Delta F_\alpha(\theta_{i+1}|\theta_i)}{f_\alpha(\theta_{i+1}|\theta_i)}\right] \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q(\theta_i|\theta_i, \theta_{i+1})dx.$$

Let $q_1 = \lim_{\alpha\to 1} q(\theta_i|\theta_i, \theta_i)$. Distortions converge to 1 along constant histories, so $q_1 = q(\theta_i|h^0)$.[40] By Assumption 3, since the static optimum is an interior solution, we have

$$s_q(\theta_i, q_1) = \frac{1-\sum_{k=i}^N \mu_k}{\mu_i} \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q_1)dx.$$

Also, letting $q_2 = \lim_{\alpha\to 1} q(\theta_i|\theta_i, \theta_{i+1})$, we have

$$s_q(\theta_i, q_2) \leq \frac{1-\sum_{k=i}^N \mu_k}{\mu_i} \cdot D. \int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q_2)dx.$$

It follows that:

$$\begin{aligned}
\Phi_q(\theta_i, q_1) &= s_q(\theta_i, q_1) - \frac{1-\sum_{k=i}^N \mu_k}{\mu_i} \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q_1)dx \\
&= 0 \geq s_q(\theta_i, q_2) - \frac{1-\sum_{k=i}^N \mu_k}{\mu_i} \cdot D \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q_2)dx \\
&> s_q(\theta_i, q_2) - \frac{1-\sum_{k=i}^N \mu_k}{\mu_i} \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q_2)dx = \Phi_q(\theta_i, q_2),
\end{aligned}$$

[40]This follows from the fact that $\lim_{\alpha\to 1}\Delta F_\alpha(\theta_i|\theta_i)/f_\alpha(\theta_i|\theta_i) = 1$.

where the strict inequality follows from $D < 1$. Since Φ is concave, we have $q_2 > q_1$.

Next, suppose $D = \lim_{\alpha\to 1} \Psi_i(f_\alpha) > 1$. Then, analogous to the steps above we show that $\lim_{\alpha\to 1} q(\theta_{i+1}|\theta_{i+1}, \theta_{i+1}) > \lim_{\alpha\to 1} q(\theta_{i+1}|\theta_{i+1}, \theta_i)$. Letting $q_3 = \lim_{\alpha\to 1} q(\theta_{i+1}|\theta_{i+1}, \theta_i)$, and $q_4 = \lim_{\alpha\to 1} q(\theta_{i+1}|\theta_{i+1}, \theta_{i+1})$, we get

$$s_q(\theta_i, q_4) = \frac{1 - \sum_{k=i+1}^{N} \mu_k}{\mu_{i+1}} \cdot \int_{\theta_{i+1}}^{\theta_i} u_{\theta,q}(x, q_4)dx$$

and,

$$s_q(\theta_i, q_3) \leq \frac{1 - \sum_{k=i+1}^{N} \mu_k}{\mu_{i+1}} \cdot D. \int_{\theta_{i+1}}^{\theta_i} u_{\theta,q}(x, q_3)dx$$

Thus, using $D > 1$, we obtain $\Phi_q(\theta_{i+1}, q_3) < \Phi_q(\theta_{i+1}, q_4)$, implying $q_4 > q_3$. ■

We now prove that (3.31) is generically satisfied. Define

$$\Gamma_i = \left\{ f_\alpha \in \Lambda \mid \lim_{\alpha\to 1} \Psi_i(f_\alpha) \neq 1 \right\}$$

We will show that Γ_i is open and dense in Λ, thereby establishing that $\Gamma = \cup_{i=1}^{N-1}\Gamma_i$ is open and dense in Λ, proving our result.

Lemma A6. *(3.31) is a generic property of* Λ.

Proof. Note that Λ is a space of functions from $[0, 1]$ to $[0, 1]^{N+1} \times [0, 1]^{N+1}$ that satisfy Assumption 2. Endow this space with the sup norm:

$$\|f\| = \sup_{\alpha\in[0,1]} \max_{i\in\{0,..,N\},\ j\in\{0,..,N\}} f_\alpha(\theta_i|\theta_j).$$

Given an $i \in (0, N)$, we proceed in two steps.

Step 1. We first prove that Γ_i is open. Assume not. Then for some $f_\alpha \in \Gamma_i$ and any ε-neighborhood $N_\varepsilon(f_\alpha)$ of f_α we can find a $f'_\alpha \in \overline{\Gamma}_i = \Lambda \setminus \Gamma_i$. It follows that there exists a sequence $(f^n_\alpha) \in \overline{\Gamma}_i$ such that $f^n_\alpha \to f_\alpha$. By definition $\lim_{\alpha\to 1} \Psi_i(f^n_\alpha) = 1$ for all n. Since $\Psi_i(f)$ is continuous in f, this implies

$$\lim_{\alpha\to 1}\Psi(f_\alpha)=\lim_{\alpha\to 1}\lim_{n\to\infty}\Psi(f_\alpha^n)=\lim_{n\to\infty}\lim_{\alpha\to 1}\Psi(f_\alpha^n)=1$$

proving that $f_\alpha\in\overline{\Gamma}_i$, a contradiction.

Step 2. Next we prove that the Γ_i is dense in Λ. To this goal we need to prove that for any $f_\alpha\in\Lambda$ and $\varepsilon>0$, there is a function f'_α such that and $\|f-f'\|<\varepsilon$ and $f'_\alpha\in\Gamma_i$. If $f_\alpha\in\Gamma_i$, then the result is immediate. Assume therefore that $f_\alpha\in\overline{\Gamma}_i$. Let $\tilde{f}$ be a constant stochastic matrix that satisfies first order stochastic dominance strictly and with $\tilde{f}(\theta_k|\theta_i)>0$ for all i, k. Fix $\overline{\varepsilon}>0$ and define

$$f'_\alpha=\varepsilon(\alpha)\tilde{f}+(1-\varepsilon(\alpha))\,f_\alpha$$

where $\varepsilon(\alpha)$ is a non negative function of α with $\varepsilon(\alpha)\leq\overline{\varepsilon}\ \forall\alpha$. Then, is easy to see that $\|f-f'\|<\overline{\varepsilon}$. Moreover,

$$\Psi_i(f_\alpha^{'})=\frac{\sum_{k=i}^{N}\left[\begin{array}{c}\frac{f_\alpha(\theta_k|\theta_{i+1})-f_\alpha(\theta_k|\theta_i)}{f_\alpha(\theta_i|\theta_{i+1})}\\+\frac{\varepsilon(\alpha)}{1-\varepsilon(\alpha)}\frac{\tilde{f}(\theta_k|\theta_{i+1})-\tilde{f}(\theta_k|\theta_i)}{f_\alpha(\theta_i|\theta_{i+1})}\end{array}\right]}{1+\frac{\varepsilon(\alpha)}{1-\varepsilon(\alpha)}\frac{\tilde{f}(\theta_i|\theta_{i+1})}{f_\alpha(\theta_i|\theta_{i+1})}}\cdot\frac{\sum_{k=i+1}^{N}\left[\begin{array}{c}\frac{f_\alpha(\theta_k|\theta_i)-f_\alpha(\theta_k|\theta_{i-1})}{f_\alpha(\theta_{i+1}|\theta_i)}\\+\frac{\varepsilon(\alpha)}{1-\varepsilon(\alpha)}\frac{\tilde{f}(\theta_k|\theta_i)-\tilde{f}(\theta_k|\theta_{i-1})}{f_\alpha(\theta_{i+1}|\theta_i)}\end{array}\right]}{1+\frac{\varepsilon(\alpha)}{1-\varepsilon(\alpha)}\frac{\tilde{f}(\theta_{i+1}|\theta_i)}{f_\alpha(\theta_{i+1}|\theta_i)}}\tag{3.32}$$

Since $f_\alpha(\theta_i|\theta_{i+1})$ and $f_\alpha(\theta_{i+1}|\theta_i)$ both converge to zero as $\alpha\to 1$, there are two cases to consider. Assume first $\lim_{\alpha\to 1}\frac{f_\alpha(\theta_i|\theta_{i+1})}{f_\alpha(\theta_{i+1}|\theta_i)}=0$. In this case choose $\varepsilon(\alpha)=f_\alpha(\theta_i|\theta_{i+1})$. We have:

$$\lim_{\alpha\to 1}\Psi_i(f_\alpha^{'})=\frac{1}{1+\tilde{f}(\theta_i|\theta_{i+1})}\left[1+M\sum_{k=i}^{N}\left(\tilde{f}(\theta_k|\theta_{i+1})-\tilde{f}(\theta_k|\theta_i)\right)\right]$$

where $M=\lim_{\alpha\to 1}\sum_{k=i+1}^{N}\left[\frac{f_\alpha(\theta_k|\theta_i)-f_\alpha(\theta_k|\theta_{i-1})}{f_\alpha(\theta_{i+1}|\theta_i)}\right]$. If $M=0$, then $\lim_{\alpha\to 1}\Psi_i(f_\alpha^{'})=1/(1+\tilde{f}(\theta_i|\theta_{i+1}))<1$; if $M=\infty$, then $\lim_{\alpha\to 1}\Psi_i(f_\alpha^{'})=\infty$. It follows that $\lim_{\alpha\to 1}\Psi_i(f_\alpha^{'})=$

1 only if M is a bounded non negative constant and:

$$M = \tilde{f}(\theta_i|\theta_{i+1}) / \sum_{k=i}^{N} \left(\tilde{f}(\theta_k|\theta_{i+1}) - \tilde{f}(\theta_k|\theta_i) \right). \tag{3.33}$$

Since $\tilde{f}$ is as a generic transition matrix that satisfies first-order stochastic dominance strictly, we have a contradiction. To see this, assume (3.33) holds and consider a matrix $\widehat{f}$ that is equal to $\tilde{f}$ except that $\widehat{f}(\theta_i|\theta_{i+1}) = \tilde{f}(\theta_i|\theta_{i+1}) - \epsilon$ and $\widehat{f}(\theta_{i+1}|\theta_{i+1}) = \tilde{f}(\theta_{i+1}|\theta_{i+1}) + \epsilon$ where $\epsilon > 0$ is a arbitrarily small. The new matrix still satisfies first order stochastic dominance strictly and it has $\widehat{f}(\theta_i|\theta_{i+1}) / \sum_{k=i}^{N} \left(\widehat{f}(\theta_k|\theta_{i+1}) - \widehat{f}(\theta_k|\theta_i) \right) < M$. We can then repeat the argument presented above choosing $\widehat{f}$ instead of $\tilde{f}$ and obtain a contradiction. In the case in which $\frac{f_\alpha(\theta_{i+1}|\theta_i)}{f_\alpha(\theta_i|\theta_{i+1})} \to 0$ we choose $\varepsilon(\alpha) = f_\alpha(\theta_{i+1}|\theta_i)$ and we proceed proving the result as in the previous case. ■

3.10.4 Proof of Proposition 4

We prove that for any μ, $|\Theta| > 2$ and a generic transition probability function, there exists an $\alpha^* < 1$, T^*, and $\delta^* < 1$ such that the first-order approach fails to be verified for any $\alpha > \alpha^*$, $T \geq T^*$ and $\delta > \delta^*$ if $\lim_{\alpha \to 1} \Psi_i(f_\alpha) \neq 1$ for some $i \in (0, N)$. Given this, the statement of the proposition follows from Lemma A6.

Note that as $\alpha \to 1$, $\Delta F_\alpha(\theta_k|\theta_k) \to 1$ and $\Delta F_\alpha(\theta_j|\theta_k) \to 0$ for $\forall k \neq j$. We have two cases to consider:

Case 1: $D = \lim_{\alpha \to 1} \Psi_i(f_\alpha) < 1$. We prove the result by showing that the FO-optimal contract violates the second period global incentive constraint $IC_{i-1,i+1}(\theta_i)$. To this end, we first make a useful observation.

Lemma A7. *$IC_{i-1,i+1}(h^{t-1})$ holds if and only if*

$$\int_{\theta_i}^{\theta_{i-1}} \begin{bmatrix} u_\theta(x, q(\theta_i|h^{t-1})) \\ -u_\theta(x, q(\theta_{i+1}|h^{t-1})) \end{bmatrix} dx + \delta \sum_{k=0}^{N} \begin{bmatrix} (f_\alpha(\theta_k|\theta_{i-1}) - f_\alpha(\theta_k|\theta_i)) \cdot \\ (U(\theta_k|h^{t-1}, \theta_i) - U(\theta_k|h^{t-1}, \theta_{i+1})) \end{bmatrix} \geq 0 \tag{3.34}$$

where $U(\theta_k|h^{t-1}, \theta_i) = U^*(\theta_k\,|h^{t-1}, \theta_i; \mathbf{q})$, as defined in (5) in the paper.

Proof. The global incentive compatibility constraint $IC_{i-1,i+1}(h^{t-1})$ can be written as:

$$U(\theta_{i-1}|h^{t-1}) - U(\theta_{i+1}|h^{t-1}) \geq u(\theta_{i-1}, q(\theta_{i+1}|h^{t-1})) - u(\theta_{i+1}, q(\theta_{i+1}|h^{t-1}))$$

$$+\delta \sum_{k=0}^{N} (f_\alpha(\theta_k|\theta_{i-1}) - f_\alpha(\theta_k|\theta_{i+1}))\, U(\theta_k|h^{t-1}, \theta_{i+1}) \tag{3.35}$$

Note that

$$U(\theta_{i-1}|h^{t-1}) - U(\theta_{i+1}|h^{t-1}) = \left(U(\theta_{i-1}|h^{t-1}) - U(\theta_i|h^{t-1})\right) + \left(U(\theta_i|h^{t-1}) - U(\theta_{i+1}|h^{t-1})\right).$$

So using $IC_{i-1,i}(h^{t-1})$ and $IC_{i,i+1}(h^{t-1})$, we have:

$$U(\theta_{i-1}|h^{t-1}) - U(\theta_{i+1}|h^{t-1}) - \begin{bmatrix} u(\theta_{i-1}, q(\theta_{i+1}|h^{t-1})) - u(\theta_{i+1}, q(\theta_{i+1}|h^{t-1})) \\ +\delta \sum_{k=0}^{N} (f_\alpha(\theta_k|\theta_{i-1}) - f_\alpha(\theta_k|\theta_{i+1}))\, U(\theta_k|h^{t}, \theta_{i+1}) \end{bmatrix} \tag{3.36}$$

$$= \begin{bmatrix} u(\theta_{i-1}, q(\theta_i|h^{t-1})) - u(\theta_i, q(\theta_i|h^{t-1})) \\ +u(\theta_i, q(\theta_{i+1}|h^{t-1})) - u(\theta_{i-1}, q(\theta_{i+1}|h^{t-1})) \end{bmatrix} +$$

$$\delta \sum_{k=0}^{N} \begin{pmatrix} (f_\alpha(\theta_k|\theta_{i-1}) - f_\alpha(\theta_k|\theta_i)) \\ \cdot (U(\theta_k|h^{t-1}, \theta_i) - U(\theta_k|h^{t-1}, \theta_{i+1})) \end{pmatrix}$$

Using (3.35) and (3.36), it follows that that $IC_{i-1,i}(h^{t-1})$ holds if and only if (3.34) holds. ■

So, $IC_{i-1,i+1}(\theta_i)$ holds if and only if

$$\int_{\theta_i}^{\theta_{i-1}} \begin{bmatrix} u_\theta(x, q(\theta_i|\theta_i)) \\ -u_\theta(x, q(\theta_{i+1}|\theta_i)) \end{bmatrix} dx + \delta \sum_{k=0}^{N} \begin{bmatrix} (f_\alpha(\theta_k|\theta_{i-1}) - f_\alpha(\theta_k|\theta_i)) \cdot \\ (U(\theta_k|\theta_i, \theta_i) - U(\theta_k|\theta_i, \theta_{i+1})) \end{bmatrix} \geq 0$$

We first note that:

$$\begin{aligned} & \sum_{k=0}^{N} (f_\alpha(\theta_k|\theta_{i-1}) - f_\alpha(\theta_k|\theta_i)) \left[U(\theta_k|\theta_i, \theta_i) - U(\theta_k|\theta_i, \theta_{i+1}) \right] \\ = \ & f_\alpha(\theta_{i-1}|\theta_{i-1}) \begin{bmatrix} U(\theta_{i-1}|\theta_i, \theta_i) \\ -U(\theta_{i-1i}|\theta_i, \theta_{i+1}) \end{bmatrix} - f_\alpha(\theta_i|\theta_i) \begin{bmatrix} U(\theta_i|\theta_i, \theta_i) \\ -U(\theta_i|\theta_i, \theta_{i+1}) \end{bmatrix} + o(\alpha) \end{aligned}$$

where $o(\alpha)$ is such that $o(\alpha) \to 0$ as $\alpha \to 1$. Either the distortions are finite in which case the first-order quantities are finite, or the distortions go to infinity in which case the non-negativity constraint binds and quantities are zero. Thus, all the quantities along non-constant histories remain finite in the limit and the associated probabilities converge to zero. Hence, $o(\alpha) \to 0$ as $\alpha \to 1$.

Next, let $\tilde{h}^t(\theta)$ be an history in which the realization is θ in every period for t periods. Using (3.4), we have:

$$\begin{aligned} U(\theta_{i-1}|\theta_i, \theta_i) - U(\theta_i|\theta_i, \theta_i) &= \sum_{t=3}^{T} (\delta \Delta F_\alpha(\theta_i\,|\theta_i\,))^{t-3} \Delta u(\theta_i|\tilde{h}^{t-1}(\theta_i)) + o(\alpha) \\ &= \sum_{t=3}^{T} \begin{bmatrix} (\delta \Delta F_\alpha(\theta_i\,|\theta_i\,))^{t-3} \\ \cdot \left[\int_{\theta_i}^{\theta_{i-1}} u_\theta(x, q(\theta_i|\tilde{h}^{t-1}(\theta_i))) dx \right] \end{bmatrix} + o(\alpha) \end{aligned}$$

Moreover:

$$U(\theta_{i-1}|\theta_i, \theta_{i+1}) - U(\theta_i|\theta_i, \theta_{i+1})$$

$$= \sum_{t=3}^{T} (\delta \Delta F_\alpha(\theta_i\,|\theta_i\,))^{t-3} \Delta u(\theta_i|\theta_i, \theta_{i+1}, \tilde{h}^{t-3}(\theta_i)) + o(\alpha)$$

$$= \sum_{t=3}^{T} \left[\begin{array}{c} (\delta \Delta F_\alpha(\theta_i\,|\theta_i\,))^{t-3} \\ \cdot \left[\int_{\theta_i}^{\theta_{i-1}} u_\theta(x, q(\theta_i|\theta_i, \theta_{i+1}, \tilde{h}^{t-3}(\theta_i)))dx \right] \end{array} \right] + o(\alpha)$$

As $\alpha \to 1$, we have by (3.8), $q(\theta_i|\tilde{h}^{t-1}(\theta_i)) \to q_1$, defined as the unique solution of:

$$s_q\left(\theta_i, q_1\right) = \frac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta q}(x, q_1)dx$$

and, $q(\theta_i|\theta_i, \theta_{i+1}, \tilde{h}^{t-3}(\theta_i)) \to q_2$, where q_2 satisfies

$$s_q(\theta_i, q_2) \leq \frac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \cdot D. \int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q_2)dx$$

As in the proof of Proposition 3, we get $q_2 > q_1$. Next,

$$\begin{aligned} & \lim_{\alpha \to 1} \sum_{k=0}^{N} \left(f_\alpha(\theta_k|\theta_{i-1}) - f_\alpha(\theta_k|\theta_i)\right)\left[U(\theta_k|\theta_i, \theta_i) - U(\theta_k|\theta_i, \theta_{i+1})\right] \\ = & \lim_{\alpha \to 1} \quad \sum_{t=3}^{T} \left(\delta \Delta F_\alpha(\theta_i\,|\theta_i\,)\right)^{t-3} \\ & \int_{\theta_i}^{\theta_{i-1}} \left[u_\theta(x, q(\theta_i|\tilde{h}^{t-1}(\theta_i))) - u_\theta(x, q(\theta_i|\theta_i, \theta_{i+1}, \tilde{h}^{t-3}(\theta_i))) \right] dx \\ = & \frac{1 - \delta^{T-2}}{1 - \delta} \int_{\theta_i}^{\theta_{i-1}} \left[u_\theta\left(x, q_1\right) - u_\theta\left(x, q_2\right)\right] dx \end{aligned}$$

Finally, as $\alpha \to 1$, $IC_{i-1,i+1}(\theta_i)$ holds only if:

$$\begin{aligned} & \lim_{\alpha \to 1} \int_{\theta_i}^{\theta_{i-1}} \left[u_\theta(x, q(\theta_i|\theta_i)) - u_\theta(x, q(\theta_{i+1}|\theta_i))\right] dx \\ \geq \; & \frac{1 - \delta^{T-2}}{1 - \delta} \int_{\theta_i}^{\theta_{i-1}} \left[u_\theta\left(x, q_2\right) - u_\theta\left(x, q_1\right)\right] dx \end{aligned} \tag{3.37}$$

The left hand side of (3.37) is clearly bounded for any δ. Since u_θ is strictly increasing in q and $q_2 > q_1$, the right hand side of (3.37) diverges to ∞ as $\delta \to 1$ and $T \to \infty$. We conclude that there exist thresholds for δ and T above which the inequality does not hold.

Case 2: $D = \lim_{\alpha\to 1} \Psi_i(f_\alpha) > 1$. We prove the result by showing that the FO-optimal contract violates the second period upward local incentive constraint $IC_{i+1,i}(\theta_{i+1})$. To this end, we first make a useful observation. Analogous to the arguments in Lemma A7 above, it is easy to show that $IC_{i+1,i}(\theta_{i+1})$ holds if and only if

$$\int_{\theta_{i+1}}^{\theta_i} \begin{bmatrix} u_\theta(x, q(\theta_i|\theta_{i+1})) \\ -u_\theta(x, q(\theta_{i+1}|\theta_{i+1})) \end{bmatrix} dx +$$

$$\delta \sum_{k=0}^{N} (f_\alpha(\theta_k|\theta_i) - f_\alpha(\theta_k|\theta_{i+1})) \cdot \begin{bmatrix} U(\theta_k|\theta_{i+1}, \theta_i) \\ -U(\theta_k|\theta_{i+1}, \theta_{i+1}) \end{bmatrix} \geq 0$$

Now,

$$\begin{aligned} & \sum_{k=0}^{N} (f_\alpha(\theta_k|\theta_i) - f_\alpha(\theta_k|\theta_{i+1})) \left[U(\theta_k|\theta_{i+1}, \theta_i) - U(\theta_k|\theta_{i+1}, \theta_{i+1})\right] \\ = \;& f_\alpha(\theta_i|\theta_i) \begin{bmatrix} U(\theta_i|\theta_{i+1}, \theta_i) \\ -U(\theta_i|\theta_{i+1}, \theta_{i+1}) \end{bmatrix} - f_\alpha(\theta_{i+1}|\theta_{i+1}) \begin{bmatrix} U(\theta_{i+1}|\theta_{i+1}, \theta_i) \\ -U(\theta_{i+1}|\theta_{i+1}, \theta_{i+1}) \end{bmatrix} + o(\alpha) \end{aligned}$$

where $o(\alpha)$ is such that $o(\alpha) \to 0$ as $\alpha \to 1$. Then, following the same steps as in case 1, we get that as $\alpha \to 1$, $IC_{i+1,i}(\theta_{i+1})$ holds as if and only if

$$\begin{aligned} & \lim_{\alpha\to 1} \int_{\theta_{i+1}}^{\theta_i} \left[u_\theta(x, q(\theta_i|\theta_{i+1})) - u_\theta(x, q(\theta_{i+1}|\theta_{i+1}))\right] dx \\ \geq \;& \frac{1-\delta^{T-2}}{1-\delta} \int_{\theta_i}^{\theta_{i-1}} \left[u_\theta\left(x, q_4\right) - u_\theta\left(x, q_3\right)\right] dx \end{aligned}$$

where $q_3 = \lim_{\alpha\to 1} q(\theta_{i+1}|\theta_{i+1}, \theta_i)$, and $q_4 = \lim_{\alpha\to 1} q(\theta_{i+1}|\theta_{i+1}, \theta_{i+1})$ are as in the proof of proposition 3 above, and $q_4 > q_3$ gives us the result. ■

3.11 Proof of Lemma 2

First, we prove a useful lemma that will be invoked in the proof of Lemma 2.

Lemma A8. *The optimal solution satisfies:* $q_L \leq \theta_L$, $q_L(L) \leq \theta_L$ *and* $q_M(L) \leq \theta_M$.

Proof. Suppose $q_L > \theta_L$. Then, decrease q_L by ε. Since it only appears on the RHS of incentive constraints and has positive coefficients, this does not violate any of the constraints. Moreover, the change in the monopolist's profit is proportional to

$$\left(\theta_L\left(q_L-\varepsilon\right)-\frac{1}{2}\left(q_L-\varepsilon\right)^2\right)-\left(\theta_L q_L-\frac{1}{2}q_L^2\right)=\left(q_L-\theta_L\right)\varepsilon-\frac{1}{2}\varepsilon^2.$$

We can choose ε small enough so that the above expression is positive, giving us a contradiction. We can similarly show that $q_L(L) \leq \theta_L$.

Next, suppose $q_M(L) > \theta_M$. Note that the second period incentive constraints after history L give

$$\Delta\theta q_L(L) \leq u_M(L) - u_L(L) \leq \Delta\theta q_M(L).$$

Without loss of generality, $IC_{ML}(L)$ can be assumed to hold as an equality. Suppose $u_M(L) - u_L(L) > \Delta\theta q_L(L)$. Then, decrease $u_M(L)$ so that $IC_{ML}(L)$ holds as an equality. This does not violate any constraints and keeps the profit of the monopolist the same.

If $IC_{LM}(L)$ holds as an equality, then we must have $q_M(L) = q_L(L) \leq \theta_L < \theta_M$, giving a contradiction. If $IC_{LM}(L)$ does not hold as an equality, then we can decrease $q_M(L)$ by ε without disturbing any of the constraints. Moreover, the change in the monopolist's profit is proportional to the following expression:

$$\left(\theta_M\left(q_M(L)-\varepsilon\right)-\frac{1}{2}\left(q_M(L)-\varepsilon\right)^2\right)-\left(\theta_M q_M(L)-\frac{1}{2}q_M(L)^2\right)$$

$$= (q_M(L) - \theta_M)\,\varepsilon - \frac{1}{2}\varepsilon^2.$$

We can choose ε small enough so that the above expression is positive, giving us a contradiction. ■

Now, we show that IR_L binds. Suppose not. Decrease U_H, U_M, U_L by the same small amount. The first period incentive compatibility constraints continue to hold and the second period constraints are unaffected. This increases the profit of the monopolist without disturbing any of the constraints, giving us a contradiction. Thus, $U_L = 0$. Next, we show that IC_{ML} binds. Suppose not. Decrease U_M by ε. Then, all the constraints are satisfied and we increase the monopolist's profit, giving us a contradiction. Using these two binding constraints we can eliminate U_L and U_M from the maximization problem. In particular, IC_{HM} can now be written as

$$U_H \geq \Delta\theta\,(q_M + q_L) + \delta\frac{3\alpha - 1}{2}\left[(u_H(M) - u_M(M)) + (u_M(L) - u_L(L))\right]$$

Also, IC_{HL} is given by

$$U_H \geq 2\Delta\theta q_L + \delta\frac{3\alpha - 1}{2}\left[u_H(L) - u_L(L)\right]$$

First, note that at least one of IC_{HM} and IC_{HL} must bind. If not, then we can decrease U_H and increase the monopolist's profit. Suppose IC_{HM} does not bind. Then, IC_{HL} must bind. Thus, we can eliminate U_H from the maximization problem. In particular, IC_{HM} can now be written as

$$\Delta\theta q_L + \delta\frac{3\alpha - 1}{2}\left[u_H(L) - u_M(L)\right] \geq \Delta\theta q_M + \delta\frac{3\alpha - 1}{2}\left[u_H(M) - u_M(M)\right] \quad (3.38)$$

Second, we claim that if IC_{ML} and IC_{HL} bind and IC_{HM} does not bind, then $IC_{HM}(L)$ binds. Suppose $u_H(L) - u_M(L) > \Delta\theta q_M(L)$. Decrease $u_H(L)$ by ε (and so U_H by $\delta(\alpha_{HH} - \alpha_{LH})\varepsilon$ and U_M by $\delta(\alpha_{MH} - \alpha_{LH})\varepsilon$), thereby, increasing the profit of the monopolist without disturbing any of the remaining constraints, giving us a contradiction. Thus, $IC_{HM}(L)$ must bind.

Using $IC_{HM}(M)$ and the binding $IC_{HM}(L)$ we can rewrite (3.38) to obtain:

$$\Delta\theta q_L + \delta\frac{3\alpha - 1}{2}\Delta\theta q_M(L) \geq \Delta\theta q_M + \delta\frac{3\alpha - 1}{2}\Delta\theta q_M(M)$$

Since IC_{HM} does not bind, it is easy to see that $q_M = \theta_M$ and $q_i(M) = \theta_i$ for any i. By Lemma A8, we have $q_L \leq \theta_L$ (and thus $q_L < \theta_M$) and $q_M(L) \leq \theta_M$. These clearly contradict the above inequality. Thus, we must have that IC_{HM} binds. ■

3.12 Proof of Lemma 4

For the reminder of the proof, it is useful to state the first-order conditions of the *WR-problem*. It is easy to see that the H type always gets the efficient quantity. After history H, moreover, quantities are always efficient, implying: $q_H = q_H(M) = q_H(L) = \theta_H$ and $q_H(H) = \theta_H, q_M(H) = \theta_M, q_L(H) = \theta_L$. The remaining first-order conditions are given by:

$$\begin{aligned}
[q_M]: &\quad \mu_M(\theta_M - q_M) - \mu_H\Delta\theta + \lambda\Delta\theta = 0 \\
[q_L]: &\quad \mu_L(\theta_L - q_L) - (\mu_H + \mu_M)\Delta\theta - \lambda\Delta\theta = 0 \\
[q_M(M)]: &\quad \mu_M\delta\alpha(\theta_M - q_M(M)) - \lambda_{HM}(M)\Delta\theta + \lambda_{LM}(M)\Delta\theta = 0 \\
[q_L(M)]: &\quad \mu_M\delta\frac{1-\alpha}{2}(\theta_L - q_L(M)) - \lambda_{ML}(M)\Delta\theta = 0 \\
[q_M(L)]: &\quad \mu_L\delta\frac{1-\alpha}{2}(\theta_M - q_M(L)) - \lambda_{HM}(L)\Delta\theta + \lambda_{LM}(L)\Delta\theta = 0
\end{aligned}$$

$$
\begin{aligned}
[q_L(L)]: &\quad \mu_L\delta\alpha\left(\theta_L - q_L(L)\right) - \lambda_{ML}(L)\Delta\theta = 0 \\
[\omega_{HM}(M)]: &\quad -\mu_H\delta\frac{3\alpha-1}{2} + \lambda\delta\frac{3\alpha-1}{2} + \lambda_{HM}(M) = 0 \\
[\omega_{ML}(M)]: &\quad \lambda_{ML}(M) - \lambda_{LM}(M) = 0 \\
[\omega_{HM}(L)]: &\quad -\lambda\delta\frac{3\alpha-1}{2} + \lambda_{HM}(L) = 0 \\
[\omega_{ML}(L)]: &\quad -\left(\mu_H + \mu_M\right)\delta\frac{3\alpha-1}{2} + \lambda_{ML}(L) - \lambda_{LM}(L) = 0
\end{aligned}
$$

We can now proceed with the proof. In the reminder of this section, we firs characterize the optimal allocation assuming $\lambda = 0$. We then derive the conditions under which the assumption of $\lambda = 0$ is admissible.

Assuming $\lambda = 0$, we have

$$
q_M = \theta_M - \frac{\mu_H}{\mu_M}\Delta\theta \quad \text{and} \quad q_L = \theta_L - \frac{\mu_H + \mu_M}{\mu_L}\Delta\theta. \tag{3.39}
$$

Clearly, $\lambda = 0$ implies $\lambda_{HM}(L) = 0$. Also, it is easy to show that $\lambda_{LM}(L) = 0$, else $q_M(L) > \theta_M$, which contradicts lemma A8. We therefore have $\lambda_{ML}(L) = \left(\mu_H + \mu_M\right)\delta\frac{3\alpha-1}{2}$, and the solution after history L is given by:

$$
q_M(L) = \theta_M \quad \text{and} \quad q_L(L) = \theta_L - \frac{\mu_H + \mu_M}{\mu_L}\frac{3\alpha-1}{2\alpha}\Delta\theta. \tag{3.40}
$$

Next, note that we must have $\lambda_{HM}(M) = \mu_H\delta\frac{3\alpha-1}{2}$ and $\lambda_{ML}(M) = \lambda_{LM}(M)$. We have two possible cases:

Case 1 (Region A1). $\lambda_{ML}(M) = \lambda_{LM}(M) = 0$. In this case:

$$
q_M(M) = \theta_M - \frac{\mu_H}{\mu_M}\frac{3\alpha-1}{2\alpha}\Delta\theta \quad \text{and} \quad q_L(M) = \theta_L \tag{3.41}
$$

For this to be a solution, we must have $\theta_M - \frac{\mu_H}{\mu_M}\frac{3\alpha-1}{2\alpha}\Delta\theta \geq \theta_L$, so $\alpha \leq \alpha_0(\mu_M)$ where

$$\alpha_0(\mu_M) = \frac{\mu_H}{3\mu_H - 2\mu_M}.$$

We conclude that for $\alpha \leq \alpha_0(\mu_M)$ the solution is given by $q_H = \theta_H$, $q_H(j) = \theta_H$, $q_j(H) = \theta_j$ for all $j = H, M, L$ in addition to (3.39)-(3.41).

Case 2 (Region A2). $\lambda_{ML}(M) = \lambda_{LM}(M) > 0$. Then, $q_M(M)$ and $q_L(M)$ are both equal to a constant q. From the first order condition with respect to $q_M(M)$ and $q_L(M)$ we have:

$$q_M(M) = q_L(M) = \frac{2\alpha}{1+\alpha}\theta_M + \frac{1-\alpha}{1+\alpha}\theta_L - \frac{\mu_H}{\mu_M}\frac{3\alpha-1}{1+\alpha}\Delta\theta. \tag{3.42}$$

We conclude that for $\alpha > \alpha_0(\mu_M)$ the solution is given by $q_H = \theta_H$, $q_H(j) = \theta_H$, $q_j(H) = \theta_j$ for all $j = H, M, L$, (3.39)-(3.40) and (3.42).

To characterize the necessary and sufficient condition for $\lambda = 0$, we need to verify that given the solution defined above, IC_{HL} is satisfied. Plugging in the values of Case 1, we obtain:

$$\theta_M - \frac{\mu_H}{\mu_M}\Delta\theta + \delta\frac{3\alpha-1}{2}\left(\theta_M - \frac{\mu_H}{\mu_M}\frac{3\alpha-1}{2\alpha}\Delta\theta\right) \geq \theta_L - \frac{\mu_H+\mu_M}{\mu_L}\Delta\theta + \delta\frac{3\alpha-1}{2}\theta_M, \tag{3.43}$$

that is,

$$\mu_M \geq \frac{\mu_L\left(1-\mu_L\right)\left(1+\frac{\delta}{\alpha}\left(\frac{3\alpha-1}{2}\right)^2\right)}{1+\mu_L\left(1+\frac{\delta}{\alpha}\left(\frac{3\alpha-1}{2}\right)^2\right)} = \mu_1^*\left(\alpha\right) \tag{3.44}$$

Plugging in the values of Case 2, we obtain:

$$\theta_M - \frac{\mu_H}{\mu_M}\Delta\theta + \delta\frac{3\alpha-1}{2}\left(\frac{2\alpha}{1+\alpha}\theta_M + \frac{1-\alpha}{1+\alpha}\theta_L - \frac{\mu_H}{\mu_M}\frac{3\alpha-1}{1+\alpha}\Delta\theta\right) \geq$$

$$\theta_L - \frac{\mu_H+\mu_M}{\mu_L}\Delta\theta + \delta\frac{3\alpha-1}{2}\theta_M, \tag{3.45}$$

that is,

$$\mu_M \geq \frac{\mu_L(1-\mu_L)\left(1+\delta\frac{(3\alpha-1)^2}{2(1+\alpha)}\right)}{1+\mu_L\left(1-\delta\frac{3\alpha-1}{1+\alpha}(1-2\alpha)\right)} = \mu_2^*(\alpha) \tag{3.46}$$

Let us define $\mu^*(\alpha) = \min\{\mu_1^*(\alpha), \mu_2^*(\alpha)\}$. We have the following result.

Lemma A9. *If α, μ_M is such that $\mu_M \geq \mu^*(\alpha)$ and $\alpha \leq \alpha_0(\mu_M)$ then the optimal contract is as described in Case 1 presented above. If $\mu \geq \mu^*(\alpha)$ and $\alpha > \alpha_0(\mu_M)$ then the optimal contract is as described in Case 2 presented above.*

Proof. We first prove that when $\alpha \leq \alpha_0(\mu_M)$, then $\mu_M \geq \mu^*(\alpha)$ implies $\mu_M \geq \mu_1^*(\alpha)$. To this end, we prove the counterpositive: when $\alpha \leq \alpha_0(\mu_M)$, $\mu_M < \mu_1^*(\alpha)$ implies $\mu_M < \mu^*(\alpha)$. Note that: 1. the left hand side of (3.43) and (3.45) are the same; 2. the right hand side of (3.43) is not larger than the right hand side of (3.45) if and only if $\frac{\mu_M}{\mu_H} \leq \frac{2\alpha}{3\alpha-1}$, that is if $\alpha \leq \alpha_0(\mu_M)$. It follows that if $\mu_M < \mu_1^*(\alpha)$, then neither (3.43) nor (3.45) hold, implying $\mu_M < \mu_2^*(\alpha)$ as well: we therefore conclude that $\mu_M < \mu^*(\alpha)$. Given this, the conditions $\mu_M \geq \mu^*(\alpha)$ *and* $\alpha \leq \alpha_0(\mu_M)$ imply the conditions $\mu_M \geq \mu_1^*(\alpha)$ *and* $\alpha \leq \alpha_0(\mu_M)$, so by the discussion presented above, the allocation described in Case 1 is an optimal solution of the *WR-problem*. By a similar argument, we can prove that when $\alpha > \alpha_0(\mu_M)$, then $\mu_M \geq \mu^*(\alpha)$ implies $\mu_M \geq \mu_2^*(\alpha)$. This implies that when we have $\mu_M \geq \mu^*(\alpha)$ *and* $\alpha > \alpha_0(\mu_M)$, then the allocation described in Case 2 is an optimal solution of the *WR-problem*. ■

Finally note that Case 1 and Case 2 described above are the only possible allocations consistent with $\lambda = 0$. So, if $\mu_M < \mu^*(\alpha)$, the Largrange multiplier of IC_{HL} must be binding.

3.13 Proof of Proposition 6

The result follows from Lemma A9. ■

3.14 Proof of Proposition 7

We first prove a useful lemma.

Lemma A10. *The optimal solution satisfies:* $q_L \leq \theta_L - \frac{\mu_H+\mu_M}{\mu_L}\Delta\theta$, $q_L(L) \leq \theta_L - \frac{\mu_H+\mu_M}{\mu_L}\frac{3\alpha-1}{2\alpha}\Delta\theta$ and $q_L(M) \leq \theta_L$.

Proof. We proceed in 3 steps.

Step 1. Suppose $q_L > \theta_L - \frac{\mu_H+\mu_M}{\mu_L}\Delta\theta$. Now, decrease q_L by ϵ. All the constraints are still satisfied. The change in the monopolist's profit is given by

$$\begin{aligned} & \mu_L\left[-\theta_L\epsilon - \frac{1}{2}\left((q_L-\epsilon)^2 - (q_L)^2\right)\right] + (\mu_H+\mu_M)\Delta\theta\epsilon \\ = \ & \mu_L\left[\left(q_L - \left(\theta_L - \frac{\mu_H+\mu_M}{\mu_L}\Delta\theta\right)\right)\epsilon - \frac{1}{2}\epsilon^2\right], \end{aligned}$$

which is greater than zero for small enough ϵ, giving us a contradiction.

Step 2. Suppose $q_L(L) > \theta_L - \frac{\mu_H+\mu_M}{\mu_L}\frac{3\alpha-1}{2\alpha}\Delta\theta$. Now, decrease $q_L(L)$ by ϵ and $\omega_{ML}(L)$ by $\Delta\theta\epsilon$. All the constraints are still satisfied. The change in the monopolist's profit is given by

$$\begin{aligned} & \mu_L\delta\alpha\left[-\theta_L\epsilon - \frac{1}{2}\left((q_L(L)-\epsilon)^2 - (q_L(L))^2\right)\right] + (\mu_H+\mu_M)\delta\frac{3\alpha-1}{2}\Delta\theta\epsilon \\ = \ & \mu_L\delta\alpha\left[\left(q_L(L) - \left(\theta_L - \frac{\mu_H+\mu_M}{\mu_L}\frac{3\alpha-1}{2\alpha}\Delta\theta\right)\right)\epsilon - \frac{1}{2}\epsilon^2\right], \end{aligned}$$

which is greater than zero for small enough ϵ, giving us a contradiction.

Step 3. Suppose $q_L(M) > \theta_L$. Now, decrease $q_L(M)$ by ϵ and $\omega_{ML}(M)$ by $\Delta\theta\epsilon$. All the constraints are still satisfied. The change in the monopolist's profit is given by

$$\mu_M\delta\frac{1-\alpha}{2}\left[-\theta_L\epsilon - \frac{1}{2}\left((q_L(M)-\epsilon)^2 - (q_L(M))^2\right)\right] =$$

$$\mu_M\delta\frac{1-\alpha}{2}\left[\left[(q_L(M)-\theta_L)\,\epsilon - \frac{1}{2}\epsilon^2\right]\right],$$

which is greater than zero for small enough ϵ, giving us a contradiction. ■

Keep in mind that $\lambda > 0 \Rightarrow \lambda_{HM}(L) > 0$. It follows from the first order condition with respect to $\omega_{HM}(L)$. Next, in order to characterize the quantities after history M, we prove a useful lemma.

Lemma A11. $\lambda > 0 \Rightarrow \lambda_{HM}(M) > 0$.

Proof. Assume by contradiction that $\lambda_{HM}(M) = 0$. Then, we must have $\lambda_{ML}(M) = \lambda_{LM}(M) = 0$. Assuming them strictly positive gives us $q_M(M) = q_L(M)$. Also, from the first order condition for $q_M(M)$, we obtain $q_M(M) > \theta_M$, implying $q_L(M) > \theta_M > \theta_L$, a contradiction to Lemma A9. Thus, $\lambda = \mu_H$ and $q_M = q_M(M) = \theta_M$.

Next, we note that if $\lambda > 0$, then $q_M(L) < \theta_M$. To see this point, consider the first-order condition with respect to $q_M(L)$. Since, $\lambda_{HM}(L) > 0$, if $\lambda_{LM}(L) = 0$ then it follows immediately that $q_M(L) < \theta_M$. If $\lambda_{LM}(L) > 0$, then $q_M(L) = q_L(L) < \theta_L < \theta_M$, where the first inequality follows from Lemma A10.

Using these facts, we can now write:

$$\Delta\theta q_M + \delta\frac{3\alpha - 1}{2}\omega_{HM}(M) = \Delta\theta \cdot \theta_M + \delta\frac{3\alpha - 1}{2}\omega_{HM}(M) \geq \Delta\theta \cdot \theta_M + \delta\frac{3\alpha - 1}{2}\Delta\theta q_M(M) \tag{3.47}$$

$$= \Delta\theta \cdot \theta_M + \delta\frac{3\alpha - 1}{2}\Delta\theta \cdot \theta_M > \Delta\theta q_L + \delta\frac{3\alpha - 1}{2}\Delta\theta q_M(L)$$
$$= \Delta\theta q_L + \delta\frac{3\alpha - 1}{2}\omega_{HM}(L).$$

The strict inequality proven in (3.47) contradicts $\lambda > 0$. Thus, we must have $\lambda_{HM}(M) > 0$ as requested. This completes the proof of Lemma A11. ■

We divide the proof of Proposition 4 into two steps. First, in Section 7.1, we assume that $IC_{LM}(L)$ is not binding and we characterize the parameter region in which this assumption is correct. This will allow us to define the regions B1 and B2 described in Proposition 7. Then, in Section 7.2, we characterize the optimal contract when $IC_{LM}(L)$ is binding, region B3.

3.14.1 Characterization of Regions B1 and B2

Let us assume $\lambda_{LM}(L) = 0$. Since $\mu_M < \mu^*(\alpha)$, we have $\lambda > 0$. From the first order conditions, we obtain:

$$
\begin{aligned}
q_M &= \theta_M - \frac{\mu_H - \lambda}{\mu_M}\Delta\theta, \;\; q_L = \theta_L - \frac{\mu_H + \mu_M + \lambda}{\mu_L}\Delta\theta && (3.48)\\
q_M(L) &= \theta_M - \frac{\lambda}{\mu_L}\frac{3\alpha - 1}{1-\alpha}\Delta\theta, \;\; q_L(L) = \theta_L - \frac{\mu_H + \mu_M}{\mu_L}\frac{3\alpha-1}{2\alpha}\Delta\theta && (3.49)
\end{aligned}
$$

Since $\lambda > 0$, we have $\lambda_{HM}(M) > 0$ and $\lambda_{HM}(L) > 0$. Thus,

$$
q_M + \delta\frac{3\alpha - 1}{2}q_M(M) = q_L + \delta\frac{3\alpha-1}{2}q_M(L) \tag{3.50}
$$

There are two relevant cases. We use λ_1 to denote λ from Case 1 and λ_2 from Case 2.

Case 1 (Region B1). $\lambda_{ML}(M) = \lambda_{LM}(M) = 0$. Then, from the first-order conditions:

$$
q_M(M) = \theta_M - \frac{\mu_H - \lambda_1}{\mu_M}\frac{3\alpha-1}{2\alpha}\Delta\theta \quad \text{and} \quad q_L(M) = \theta_L \tag{3.51}
$$

Substituting, the values from (3.48)-(3.49) and (3.51) in equation (3.50) we obtain:

$$
\frac{1+\lambda_1}{\mu_L} + \delta\frac{3\alpha-1}{2}\frac{\lambda_1}{\mu_L}\frac{3\alpha-1}{1-\alpha} = \frac{\mu_H - \lambda_1}{\mu_M} + \delta\frac{3\alpha-1}{2}\frac{\mu_H-\lambda_1}{\mu_M}\frac{3\alpha-1}{2\alpha} \tag{3.52}
$$

which gives:

$$
\lambda_1 = \lambda_1(\alpha) = \frac{\frac{\mu_H}{\mu_M}\left(1 + \delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{2\alpha}\right) - \frac{1}{\mu_L}}{\frac{1}{\mu_M}\left(1+\delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{2\alpha}\right) + \frac{1}{\mu_L}\left(1+\delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{1-\alpha}\right)} \tag{3.53}
$$

Clearly, for this case to be valid, we must justify the assumption that $\lambda_{ML}(M) =$

$\lambda_{LM}(M) = 0$. A necessary and sufficient condition for this is $q_M(M) \geq q_L(M)$. Given (3.51), this condition can be rewritten as: $\frac{\mu_H - \lambda_1}{\mu_M}\frac{3\alpha-1}{2\alpha} \leq 1$, where λ_1 is given by (3.53). This condition is implied by:

$$\mu_M \geq \frac{1 + (1-\mu_L)b_0(\alpha) - \mu_L c_0(\alpha) a_0(\alpha)}{b_0(\alpha)\left(1 + c_0(\alpha)\right)} = \mu_0(\alpha),$$

where

$$a_0(\alpha) = 1 + \delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{2\alpha},\ b_0(\alpha) = 1 + \delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{1-\alpha},\ \text{and } c_0(\alpha) = \frac{2\alpha}{3\alpha-1}$$

It follows that (under the assumption that $\lambda_{LM}(L) = 0$) the solution is given by (3.48)-(3.49), (3.51) and (3.53) when when $\mu_M \geq \mu_0(\alpha)$.

Case 2 (Region B2). For $\mu_M < \mu_0(\alpha)$ we must have $\lambda_{ML}(M) = \lambda_{LM}(M) > 0$. In this case, we must have:

$$q_M(M) = q_L(M) = \frac{2\alpha}{1+\alpha}\theta_M + \frac{1-\alpha}{1+\alpha}\theta_L - \frac{\mu_H - \lambda_2}{\mu_M}\frac{3\alpha-1}{1+\alpha}\Delta\theta \tag{3.54}$$

Substituting $q_M(M)$ and $q_M(L)$ equation (3.50) we obtain:

$$\frac{1+\lambda_2}{\mu_L} + \delta\frac{3\alpha-1}{2}\left(\frac{\lambda_2}{\mu_L}\frac{3\alpha-1}{1-\alpha} - \frac{1-\alpha}{1+\alpha}\right) = \frac{\mu_H - \lambda_2}{\mu_M} + \delta\frac{3\alpha-1}{2}\frac{\mu_H-\lambda_2}{\mu_M}\frac{3\alpha-1}{1+\alpha} \tag{3.55}$$

which gives

$$\lambda_2 = \frac{\frac{\mu_H}{\mu_M}\left(1 + \delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{1+\alpha}\right) - \left(\frac{1}{\mu_L} - \delta\frac{3\alpha-1}{2}\frac{1-\alpha}{1+\alpha}\right)}{\frac{1}{\mu_M}\left(1 + \delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{1+\alpha}\right) + \frac{1}{\mu_L}\left(1 + \delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{1-\alpha}\right)} \tag{3.56}$$

It follows that (under the assumption that $\lambda_{LM}(L) = 0$) the solution is given by (3.48)-(3.49), (3.54) and (3.56) when $\mu_M < \mu_0(\alpha)$.

We now complete the analysis of this section by characterizing the conditions under which we can ignore the $IC_{LM}(L)$ constraint and so $\lambda_{LM}(L) = 0$. It is easy to see that $IC_{LM}(L)$ is satisfied if and only if $q_M(L) \geq q_L(L)$. We have $q_M(L) \geq q_L(L)$ if and only if:

$$\lambda_i \leq \left(\frac{1}{\mu_L}\frac{3\alpha-1}{1-\alpha}\right)^{-1}\left(1+\frac{1-\mu_L}{\mu_L}\frac{3\alpha-1}{2\alpha}\right) \tag{3.57}$$

Thus, for Case 1 we have,

$$\frac{\frac{\mu_H}{\mu_M}\left(1+\delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{2\alpha}\right)-\frac{1}{\mu_L}}{\frac{1}{\mu_M}\left(1+\delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{2\alpha}\right)+\frac{1}{\mu_L}\left(1+\delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{1-\alpha}\right)} \leq \left(\frac{1}{\mu_L}\frac{3\alpha-1}{1-\alpha}\right)^{-1}\left(1+\frac{1-\mu_L}{\mu_L}\frac{3\alpha-1}{2\alpha}\right)$$

Define

$$a_1(\alpha,\mu_L) = 1+\delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{2\alpha},\ b_1(\alpha,\mu_L) = 1+\delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{1-\alpha},\ \text{and}$$

$$c_1(\alpha,\mu_L) = \left(\frac{1}{\mu_L}\frac{3\alpha-1}{1-\alpha}\right)^{-1}\left(1+\frac{1-\mu_L}{\mu_L}\frac{3\alpha-1}{2\alpha}\right)$$

We can then write the previous inequality as:

$$\mu_M \geq \frac{\mu_L a_1(\alpha,\mu_L)\left[1-\mu_L-c_1(\alpha,\mu_L)\right]}{1+a_1(\alpha,\mu_L)\mu_L+b_1(\alpha,\mu_L)c_1(\alpha,\mu_L)} = \mu_1^{**}(\alpha)$$

Next, for Case 2, we have $q_M(L) \geq q_L(L)$ iff,

$$\frac{\frac{\mu_H}{\mu_M}\left(1+\delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{1+\alpha}\right)-\left(\frac{1}{\mu_L}-\delta\frac{3\alpha-1}{2}\frac{1-\alpha}{1+\alpha}\right)}{\frac{1}{\mu_M}\left(1+\delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{1+\alpha}\right)+\frac{1}{\mu_L}\left(1+\delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{1-\alpha}\right)} \leq \left(\frac{1}{\mu_L}\frac{3\alpha-1}{1-\alpha}\right)^{-1}\left(1+\frac{1-\mu_L}{\mu_L}\frac{3\alpha-1}{2\alpha}\right)$$

Define

$$a_2(\alpha,\mu_L) = 1+\delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{1+\alpha},\ b_2(\alpha,\mu_L) = \frac{1}{\mu_L}-\delta\frac{3\alpha-1}{2}\frac{1-\alpha}{1+\alpha}$$

$$c_2(\alpha,\mu_L) = \left(\frac{1}{\mu_L}\frac{3\alpha-1}{1-\alpha}\right)^{-1}\left(1+\frac{1-\mu_L}{\mu_L}\frac{3\alpha-1}{2\alpha}\right),$$

$$d_2(\alpha,\mu_L) = 1+\delta\frac{3\alpha-1}{2}\frac{3\alpha-1}{1-\alpha}$$

Rearranging, we obtain:

$$\mu_M \geq \frac{\mu_L a_2(\alpha, \mu_L)[1-\mu_L - c_2(\alpha, \mu_L)]}{\mu_L (a_2(\alpha, \mu_L) + b_2(\alpha, \mu_L)) + b_2(\alpha, \mu_L) c_2(\alpha, \mu_L)} = \mu_2^{**}(\alpha)$$

Let us define $\mu^{**}(\alpha) = \min\{\mu^*(\alpha), \mu_1^{**}(\alpha), \mu_2^{**}(\alpha)\}$. We have:

Lemma A12. *If $\mu_M \in [\mu^{**}(\alpha), \mu^*(\alpha)]$ and $\mu_M \geq \mu_0(\alpha)$, then the solution of the WR-problem is given by the solution in Case 1 presented above. If $\mu_M \in [\mu^{**}(\alpha), \mu^*(\alpha)]$ and $\mu_M < \mu_0(\alpha)$, then the solution of the WR-problem is given by the solution in Case 2 presented above.*

Proof. We first show that if $\mu_M \in [\mu^{**}(\alpha), \mu^*(\alpha)]$ *and* $\mu_M \geq \mu_0(\alpha)$, then $\mu_M \in [\mu_1^{**}(\alpha), \mu^*(\alpha)]$ and $\mu_M \geq \mu_0(\alpha)$. This implies that the solution is given by Case 1. Assume $\mu_M < \mu_1^{**}(\alpha)$. In this case, (3.57) does not hold with λ_1. This implies that (3.57) does not hold with λ_2 as well if $\lambda_2 \geq \lambda_1$. Subtracting equation (3.55) from equation (3.52), we get

$$(\lambda_1 - \lambda_2)\left[\frac{1}{\mu_L} + \frac{1}{\mu_M} + \delta\frac{3\alpha-1}{2}\left(\frac{1}{\mu_L}\frac{3\alpha-1}{1-\alpha} + \frac{1}{\mu_M}\frac{3\alpha-1}{1+\alpha}\right)\right] =$$

$$\delta\frac{3\alpha-1}{2}\frac{1-\alpha}{1+\alpha}\left[\frac{\mu_H - \lambda_1}{\mu_M}\frac{3\alpha-1}{2\alpha} - 1\right] \tag{3.58}$$

So, we have that $\lambda_2 \geq \lambda_1$ if:

$$\frac{\mu_H - \lambda_1}{\mu_M}\frac{3\alpha-1}{2\alpha} - 1 \leq 0,$$

which is implied by $\mu_M \geq \mu_0(\alpha)$. It follows that if $\mu_M < \mu_1^{**}(\alpha)$, then $\mu_M < \mu^{**}(\alpha)$, a contradiction. We conclude that it must be $\mu_M \geq \mu_1^{**}(\alpha)$.

We now show that if $\mu_M \in [\mu^{**}(\alpha), \mu^*(\alpha)]$ and $\mu_M < \mu_0(\alpha)$, then $\mu_M \in [\mu_2^{**}(\alpha), \mu^*(\alpha)]$ *and* $\mu_M < \mu_0(\alpha)$. This implies that the solution is given by Case 2. Assume $\mu_M < \mu_2^{**}(\alpha)$. In this case, (3.57) does not hold with λ_2 . This implies that (3.57) does not hold with λ_1 as well if $\lambda_1 \geq \lambda_2$. From (3.58) we have that this always true if $\mu_M < \mu_0(\alpha)$. It follows that if $\mu_M < \mu_2^{**}(\alpha)$, then $\mu_M < \mu^{**}(\alpha)$, a contradiction. We conclude that it must be $\mu_M \geq \mu_2^{**}(\alpha)$. ■

3.14.2 Characterization of Region B3

Finally, we characterize the contract when $\mu_M < \mu^{**}(\alpha)$ and so both $\lambda > 0$ and $\lambda_{LM}(L) > 0$. This is region B3. In this case:

$$q_M = \theta_M - \frac{\mu_H - \lambda}{\mu_M}\Delta\theta \quad \text{and} \quad q_L = \theta_L - \frac{\mu_H + \mu_M + \lambda}{\mu_L}\Delta\theta \tag{3.59}$$

We also have that $\lambda_{LM}(L) > 0$ implies $q_M(L) = q_L(L)$, so:

$$q_M(L) = q_L(L) = \frac{1-\alpha}{1+\alpha}\theta_M + \frac{2\alpha}{1+\alpha}\theta_L - \frac{\mu_H + \mu_M + \lambda}{\mu_L}\frac{3\alpha - 1}{1+\alpha}\Delta\theta \tag{3.60}$$

From Lemma A9, we have $q_L(L) \leq \theta_L - \frac{\mu_H+\mu_M}{\mu_L}\frac{3\alpha-1}{2\alpha}\Delta\theta$. Also, when $\lambda_{LM}(L) > 0$, the above inequality is strict. Thus, substituting the optimal value of $q_L(L)$, we obtain:

$$1 - \frac{\lambda}{\mu_L}\frac{3\alpha - 1}{1-\alpha} + \frac{\mu_H + \mu_M}{\mu_L}\frac{3\alpha - 1}{2\alpha} < 0 \tag{3.61}$$

Note that as $\lambda_{LM}(L)$ converges to zero, (3.61) is the exact violation of $\mu_M \geq \mu^{**}(\alpha)$, that is, inequality (3.57).

To characterize the quantities after history M, we now show that $\lambda_{ML}(M) = \lambda_{LM}(M) > 0$.

Lemma A13. $\lambda, \lambda_{LM}(L) > 0 \Rightarrow \lambda_{ML}(M) = \lambda_{LM}(L) > 0$.

Proof. Suppose $\lambda_{ML}(M) = \lambda_{LM}(M) = 0$. Then,

$$q_M(M) = \theta_M - \frac{\mu_H - \lambda}{\mu_M}\frac{3\alpha - 1}{2\alpha}\Delta\theta \quad \text{and} \quad q_L(M) = \theta_L.$$

From $\theta_M - \frac{\mu_H-\lambda}{\mu_M}\frac{3\alpha-1}{2\alpha}\Delta \geq \theta_L$, we have:

$$\frac{2\alpha}{3\alpha-1} - \frac{\mu_H-\lambda}{\mu_M} \geq 0. \tag{3.62}$$

Since $\lambda, \lambda_{LM}(L) > 0$, using $q_M(M) \geq q_L(M) = \theta_L > q_L(L) = q_M(L)$, we get $q_L > q_M$. This implies

$$\left(1 - \frac{\mu_H-\lambda}{\mu_M} + \frac{\mu_H+\mu_M+\lambda}{\mu_L}\right) < 0.$$

Using equation (3.62). we get

$$\frac{\mu_H+\mu_M+\lambda}{\mu_L}\frac{3\alpha-1}{1-\alpha} < 1. \tag{3.63}$$

Now, inequality (3.61) can be written as

$$1 < \frac{\mu_H+\mu_M+\lambda}{\mu_L}\frac{3\alpha-1}{1-\alpha} - \frac{\mu_H+\mu_M}{\mu_L}(3\alpha-1)\left(\frac{1}{1-\alpha}+\frac{1}{2\alpha}\right)$$

$$= \frac{\mu_H+\mu_M+\lambda}{\mu_L}\frac{3\alpha-1}{1-\alpha} - \frac{\mu_H+\mu_M}{\mu_L}\frac{3\alpha-1}{2\alpha}\frac{1+\alpha}{2\alpha}$$

which contradicts condition (3.63). ■

It follows that

$$q_M(M) = q_L(M) = \frac{2\alpha}{1+\alpha}\theta_M + \frac{1-\alpha}{1+\alpha}\theta_L - \frac{\mu_H-\lambda}{\mu_M}\frac{3\alpha-1}{1+\alpha}\Delta\theta$$

Finally, substituting the optimal values in IC_{HL} as equality, we obtain:

$$\left(1 - \frac{\mu_H-\lambda}{\mu_M} + \frac{\mu_H+\mu_M+\lambda}{\mu_L}\right) = 0 \tag{3.64}$$

that implies $q_M = q_L$. Note that equation (3.64) gives the value of λ, which uniquely defines the solution at the optimum. In particular note that type M and L are treated as *one*, that is,

$$q_M = q_L \quad \text{and} \quad q_M(M) = q_L(M) = q_M(L) = q_L(L) \tag{3.65}$$

We conclude that the solution of the *WR-problem* in region B3 ($\mu_M < \mu^{**}(\alpha)$) is given by: (3.59),(3.60), (3.65) and (3.64). ■

Table 1 summarizes the solution of the *WR-problem* describing the optimal allocation for each possible case.

3.14.3 Proof of Proposition 8

We proceed in two steps.

Step 1. We say that a quantity $q\left(\theta_i \left| h^{t-1}\right.\right)$ is distorted downward (respectively, upward) if $q\left(\theta_i \left| h^{t-1}\right.\right) \leq \theta_i$ (respectively, $q\left(\theta_i \left| h^{t-1}\right.\right) > \theta_i$). We first show that in the optimal monotonic contract distortions are all downward. Consider the constraint set of (3.20), as described by $\mathcal{M}$. Define,

$$\Gamma\left(h^{t-1}\right) = \left\{ \begin{array}{c} \widehat{h}^{t-1} | \exists k \leq t-1 \text{ s.t. given } h_k^{t-1} = \theta_l \text{ for some } l = 0,1,...,N-1; \\ \text{we have } \widehat{h}_k^{t-1} = \theta_{l+1} \text{ and } h_j^{t-1} = \widehat{h}_j^{t-1} \ \forall j \neq k \end{array} \right\}$$

Thus, $\Gamma\left(h^{t-1}\right)$ is the set of histories that differ from h^{t-1} only once: the type in period k is replaced by the contiguous lower type. It is easy to see that a contract is monotonic if and only if for any history h^{t-1}: 1. $q\left(\theta_i|h^{t-1}\right) \geq q\left(\theta_{i+1}|h^{t-1}\right)$ for all $i < N$; and, 2. $q\left(\theta_i|h^{t-1}\right) \geq q\left(\theta_i|\widehat{h}^{t-1}\right)$ for all i and for all $\widehat{h}^{t-1} \in \Gamma\left(h^{t-1}\right)$.

Next, we introduce the following complete order on the set of all histories at time t. For any two histories h^{t-1} and $\widehat{h}^{t-1}$, let $\tau^*\left(h^{t-1}, \widehat{h}^{t-1}\right)$ be the first period in which they diverge: $\tau^*\left(h^{t-1}, \widehat{h}^{t-1}\right) = \min_j \left\{0 \leq j \leq t-1 \text{ s.t. } h_j^{t-1} \neq \widehat{h}_j^{t-1}\right\}$, with $\tau^*\left(h^{t-1}, \widehat{h}^{t-1}\right) = t-1$ if $h^{t-1} = \widehat{h}^{t-1}$. We say that $h^{t-1} \succeq^* \widehat{h}^{t-1}$ if $h^{t-1}_{\tau^*\left(h^{t-1}, \widehat{h}^{t-1}\right)} \geq$

$\widehat{h}^{t-1}_{\tau^*\left(h^{t-1},\widehat{h}^{t-1}\right)}$, i.e., if it is higher at the first point of divergence. It is easy to verify that the order $\succeq^*$ is complete, so without loss we can order the histories at time t from largest ($\overline{h}^{t-1}$) to smallest ($\underline{h}^{t-1}$), where the largest (smallest) history has all realizations equal to θ_0 (θ_N). Also, note that, $h^{t-1} \succeq^* \widehat{h}^{t-1}$ for all $\widehat{h}^{t-1} \in \Gamma\left(h^{t-1}\right)$.

Consider period t, and the smallest history of length $t-1$ (denoted, $\underline{h}^{t-1}$), in which all the realizations are θ_N. It is immediate to see that $q\left(\theta_N|\underline{h}^{t-1}\right)$ can not be distorted upward. To see this note that $q\left(\theta_N\left|\underline{h}^{t-1}\right.\right)$ is on the left hand side of no constraint.[41] If it were distorted upward, then a marginal decrease in $q\left(\theta_N|\underline{h}^{t-1}\right)$ would relax all constraints and increase surplus. Now, consider $q\left(\theta_{N-1}|\underline{h}^{t-1}\right)$: this quantity appears on the left hand side of only one constraint; $q\left(\theta_{N-1}|\underline{h}^{t-1}\right) \geq q\left(\theta_N|\underline{h}^{t-1}\right)$. If this constraint is not binding, then by the argument presented above, $q\left(\theta_{N-1}|\underline{h}^{t-1}\right) \leq \theta_{N-1}$. Assume it is binding. In this case $q\left(\theta_{N-1}|\underline{h}^{t-1}\right) = q\left(\theta_N|\underline{h}^{t-1}\right) \leq \theta_N \leq \theta_{N-1}$. Proceeding inductively with a similar argument, we can prove that $q\left(\theta_i\left|\underline{h}^{t-1}\right.\right) \leq \theta_i$ for all i.

Note that the case for first period quantities, when the history is just the empty set, is already covered by the above paragraph. Thus, now we consider $t \geq 2$. Assume, as an induction step, that there is a history $\widehat{h}^{t-1}$, where $\widehat{h}^{t-1} \succeq^* \underline{h}^{t-1}$, such that $\widehat{h}^{t-1} \succeq^* h^{t-1} \succeq^* \underline{h}^{t-1}$ implies $q\left(\theta_i|h^{t-1}\right) \leq \theta_i$ for all i. Let us also introduce a useful definition. For any h^{t-1} with $\overline{h}^{t-1} \succeq^* h^{t-1}, h^{t-1} \neq \overline{h}^{t-1}$ and $t \geq 2$, define $[h^{t-1}]^+$ to be the smallest t-period history larger than h^{t-1} according to the order $\succeq^*$in the following inductive way. If $t=2$, then $[h^{t-1}]^+ = \left\{\kappa_{t-1}(h^{t-1}), h^{t-1}_{t-1} + \Delta\theta\right\}$; if $t > 2$ then:

$$\left[h^{t-1}\right]^+ = \begin{cases} \left(\kappa_{t-1}(h^{t-1}), h^{t-1}_{t-1} + \Delta\theta\right), & \text{if } h^{t-1}_{t-1} < \theta_0 \\ \left(\left[\kappa_{t-1}(h^{t-1})\right]^+, \theta_N\right), & \text{if } h^{t-1}_{t-1} = \theta_0 \end{cases},$$

where κ_s projects the first s elements of a vector.[42] We intend to show that

[41]We say that a quantity is on the left hand side of a given constraint if in that constraint it must be larger than some other quantity.

[42]Recollect that h^{t-1} is a vector of length t: $h^{t-1} = (h^{t-1}_0, h^{t-1}_1, ..., h^{t-1}_{t-1})$, where

$q\left(\theta_i|\left[\widehat{h}^{t-1}\right]^+\right) \leq \theta_i$ for all i. Now, $q\left(\theta_N|\left[\widehat{h}^{t-1}\right]^+\right)$ appears on the left hand side in the following constraints: $q\left(\theta_N|\left[\widehat{h}^{t-1}\right]^+\right) \geq q\left(\theta_N|\widetilde{h}^{t-1}\right)$ for all $\widetilde{h}^{t-1} \in \Gamma\left(\left[\widehat{h}^{t-1}\right]^+\right)$. If none of these constraints bind, then as before, we have the desired inequality. Suppose at least one of them binds. Clearly, by the definition of $\left[\widehat{h}^{t-1}\right]^+$, we have $\widehat{h}^{t-1} \succeq^* \widetilde{h}^{t-1}$ for all $\widetilde{h}^{t-1} \in \Gamma\left(\left[\widehat{h}^{t-1}\right]^+\right)$. Thus, by the induction hypothesis $q\left(\theta_N|\widetilde{h}^{t-1}\right) \leq \theta_N$ for all $\widetilde{h}^{t-1} \in \Gamma\left(\left[\widehat{h}^{t-1}\right]^+\right)$. Since the inequality constraint binds for some $\widetilde{h}^{t-1}$, we have $q\left(\theta_N|\left[\widehat{h}^{t-1}\right]^+\right) = q\left(\theta_N|\widetilde{h}^{t-1}\right) \leq \theta_N$.

Next, consider $q\left(\theta_{N-1}|\left[\widehat{h}^{t-1}\right]^+\right)$. It appears on the left hand side in the following constraints: $q\left(\theta_{N-1}|\left[\widehat{h}^{t-1}\right]^+\right) \geq q\left(\theta_N|\left[\widehat{h}^{t-1}\right]^+\right)$ and $q\left(\theta_{N-1}|\left[\widehat{h}^{t-1}\right]^+\right) \geq q\left(\theta_{N-1}|\widetilde{h}^{t-1}\right)$ for all $\widetilde{h}^{t-1} \in \Gamma\left(\left[\widehat{h}^{t-1}\right]^+\right)$. If none of these constraints bind, then as before, we have the desired inequality. If the first one binds then, $q\left(\theta_{N-1}|\left[\widehat{h}^{t-1}\right]^+\right) \leq \theta_N < \theta_{N-1}$. If any of the latter one binds, then invoking the induction hypothesis, as argued in the case above, we have the desired inequality. Proceeding inductively, we can show $q\left(\theta_i|h^{t-1}\right) \leq \theta_i$ for all i and h^{t-1}.

Step 2. We now prove that the allocation is asymptotially efficient. Consider problem (3.20). From this problem eliminate the constraint $q\left(\theta_0|h^0\right) \geq q\left(\theta_1|h^0\right)$ and all the monotonicity constraints that involve quantities following an history in which the agents reports to be a type θ_0. It is easy to see that in this problem the quantities offered after the agent reports (or has reported) to be θ_0 are efficient: $q\left(\theta_i|h^{t-1}\right) = \theta_i$ for $i = 0$ and/or $\forall h^{t-1} \in \overline{H}^{t-1}, t \geq 2$, where $\overline{H}^{t-1} = \{h^{t-1}\,|\exists\tau \leq t-1 \text{ s.t. } h_\tau^{t-1} = \theta_0\}$. Following the same approach as in Step 1, it can be shown that the solution of this relaxed problem is monotonic and so it coincides with the optimal monotonic contract. Since the probability of the event in which no type realization in t periods is equal to θ_0 converges to zero as $t \to \infty$, this solution is, is asymptotically efficient, and so the

$h_0^{t-1} = \emptyset$. So, $\kappa_{t-1}(h^{t-1}) = (h_0^{t-1}, ..., h_{t-2}^{t-1})$.

optimal monotonic contract. ■

3.14.4 Proof of Proposition 9

We prove that for any given T, the optimal monotonic contract converges in probability to the optimal contract. Let $\Pi^s(\alpha)$, $\Pi^m(\alpha)$ and $\Pi^{**}(\alpha)$ be the expected profits obtained by the seller from, respectively, the repetition of the optimal static contract, the optimal monotonic contract and the optimal contract when the Markov matrix is α. Because the repetition of the optimal static contract is a monotonic dynamic contract, we must have $\Pi^m(\alpha) \in [\Pi^s(\alpha), \Pi^{**}(\alpha)]$. Now note that when types are constant and $\alpha = I$, it is well known that the repetition in every period of the optimal static contract is optimal.[43] Since $\Pi^m(\alpha), \Pi^s(\alpha)$ and $\Pi^{**}(\alpha)$ are continuous in α by the theorem of the maximum[44], we must have that for any sequence $\alpha_n \to I$ and $\varepsilon > 0$ there must be a n' such that for $n > n'$, we have $|\Pi^m(\alpha_n) - \Pi^{**}(\alpha_n)| \leq |\Pi^s(\alpha_n) - \Pi^{**}(\alpha_n)| < \varepsilon$. It is immediate to see that the fact that $\Pi^m(\alpha_n)$ converges to $\Pi^{**}(\alpha_n)$ and that by Proposition 8 quantities are bounded imply that the optimal monotonic contract must converge to a contract that maximizes profit in probability. ■

[43] This result can be easily deduced studying problem (3.20). To see it, note that the repetition of the static contract is incentive compatible and individually rational. Then note that when $\Lambda = I$, the first order optimal contract coincides with the static optimal contract along the histories in which the agent reports always the same type. Since the other histories have probability zero, the profit from the repetition of the static contract is the same of the profit from the FO-optimal contract. Since the FO-optimal contract yields a profit not inferior to the optimal contract, the result is proven.

[44] In order to apply the theorem of the maximum the space of quantities must be compact. It is clearly bounded below by zero. Also, Proposition 8 shows that it is bounded above by the efficient quantities. Hence, there is no loss of generality in assuming that set of quantities is contained in the interval $[0, \theta_0\ ]$.

3.14.5 Proof of Proposition 10

The fact that $\lim_{\delta\to 1}\lim_{\alpha\to I}\pi_m(\alpha,\delta)=\lim_{\delta\to 1}\lim_{\alpha\to I}\pi^*(\alpha,\delta)$ follows immediately from the fact that for any δ, $\lim_{\alpha\to I}\pi_m(\alpha,\delta)=\lim_{\alpha\to I}\pi^*(\alpha,\delta)$. We now prove the remaining equality. Let $S(\alpha,\delta,T)$ be the total expected surplus generated in the efficient contract, $U^*(\alpha,\delta,T)$ be the agent's expected surplus obtained with the efficient contract and $U_i^*(\alpha,\delta,T)$ be the agent's surplus obtained with the efficient contract conditional on being type i at $t=1$. Define also $s(\alpha,\delta,T)=(1-\delta)\,S(\alpha,\delta,T)$, $u^*(\alpha,\delta,T)=(1-\delta)\,U^*(\alpha,\delta,T)$ and $u_i^*(\alpha,\delta,T)=(1-\delta)\,U_i^*(\alpha,\delta,T)$; and $s(\alpha,\delta)=\lim_{T\to\infty}s(\alpha,\delta,T)$, $u^*(\alpha,\delta)=\lim_{T\to\infty}u^*(\alpha,\delta,T)$ and $u_i^*(\alpha,\delta)=\lim_{T\to\infty}u_i^*(\alpha,\delta,T)$. Profits $\pi_m(\alpha,\delta)$ must be larger or equal to the profits obtained by offering the efficient quantity and charging a fixed per period price equal to $u_N^*(\alpha,\delta)$, since this is an incentive compatible monotonic contract. Note that since types follow an irreducible Markov process, their distribution converges to a stationary distribution that is independent from the realization at $t=1$. It follows that, for all α, $\lim_{\delta\to 1}u_i^*(\alpha,\delta)=\lim_{\delta\to 1}u^*(\alpha,\delta)$ and so the per period profits in this contract converge to $s(\alpha,\delta)$, implying that, for all α, $\lim_{\delta\to 1}\pi_m(\alpha,\delta)=\lim_{\delta\to 1}s(\alpha,\delta)$. Similarly we can show that for all α, $\lim_{\delta\to 1}\pi^*(\alpha,\delta)=\lim_{\delta\to 1}s(\alpha,\delta)$. It follows that: $\lim_{\alpha\to I}\lim_{\delta\to 1}\pi_m(\alpha,\delta)=\lim_{\alpha\to I}\lim_{\delta\to 1}s(\alpha,\delta)=\lim_{\alpha\to I}\lim_{\delta\to 1}\pi^*(\alpha,\delta)$. This proves the result. ■

3.15 Proof of Lemma 3

We prove the lemma as follows. Let $\mathbf{U}=U(h^t)$ be the vector of expected utilities, mapping an history h^t to the corresponding agent's expected utility. First, we construct a vector of utilities $\mathbf{U}$ using the solution of the *WR-problem*, $\langle\omega,\mathbf{q}\rangle$. We then show that the solution $\langle\mathbf{U},\mathbf{q}\rangle$ satisfies all the constraints of the seller's profit maximization problem and it maximizes profits. We proceed in two steps:

Step 1. We set $u_L(M)$, $u_L(L)$, $u_L(H)$ all equal to zero. We also define:

$$u_M(M) = \omega_{ML}(M), u_M(L) = \omega_{ML}(L), u_M(H) = \Delta\theta q_L(H)$$
$$u_H(M) = \omega_{ML}(M) + \omega_{HM}(M), u_H(L) = \omega_{ML}(L) + \omega_{HM}(L), u_H(H)$$
$$= \Delta\theta \left(q_L(H) + q_M(H)\right)$$

Since IR_L, IC_{ML} and IC_{HM} hold as an equality, we must have:

$$\begin{aligned} U_L &= 0, \\ U_M &= \Delta\theta q_L + \delta\frac{3\alpha-1}{2}\omega_{ML}(L), \text{ and} \\ U_H &= U_M + \Delta\theta q_M + \delta\frac{3\alpha-1}{2}\omega_{HM}(M). \end{aligned}$$

Step 2. We now show that $\langle \mathbf{U}, \mathbf{q} \rangle$ satisfies all the constraints of the profit maximizing problem. By construction it is immediate that $\langle \mathbf{U}, \mathbf{q} \rangle$ satisfies all the constraints in the *WR-problem.* It remains to be shown that it also satisfies the other constraints,

$$IR_H, IR_M, IC_{MH}, IC_{LM}, IC_{LH}, \tag{3.66}$$
$$IC_{HM}(H), IC_{ML}(H), IR_L(H), IR_L(M), IR_L(L)$$
$$IC_{MH}(H), IC_{LM}(H), IC_{LH}(H), IC_{HL}(H)), IC_{MH}(M),$$
$$IC_{LH}(M), IC_{HL}(M), IC_{MH}(L), IC_{LH}(L), IC_{HL}(L).$$

First, we show that IR_M is satisfied. From IC_{ML} we have

$$\begin{aligned} U_M &= U_L + \Delta\theta q_L + \delta\frac{3\alpha-1}{2}\left[u_M(L) - u_L(L)\right] \\ &= \Delta\theta q_L + \delta\frac{3\alpha-1}{2}\left[u_M(L) - u_L(L)\right] \qquad \text{[Using } IR_L\text{]} \\ &\geq \Delta\theta q_L + \delta\frac{3\alpha-1}{2}\Delta q_L(L) > 0 \qquad \text{[Using } IC_{ML}(L)\text{]} \end{aligned}$$

Similarly, we can show that IR_H is satisfied. To prove the remaining constraints we need the following properties of the solution of the WR-problem.

Lemma A14. *For all parameter configurations, in the solution to the WR-problem we have: 1.* $q_i(H) = \theta_i$ *for* $i = M, L, H$*,* $q_M(M) < \theta_M, q_L(M) \leq \theta_L$*, and* $q_L(M) \geq q_L(L)$ *2.* $\omega_{HM}(M) = \Delta\theta q_M(M)$ *and, without loss of generality,* $\omega_{ML}(M) = \Delta\theta q_L(M)$*,* $\omega_{HM}(L) = \Delta\theta q_M(L)$*; 3. quantities at* $t = 2$ *are nondecreasing in type after any history; 4.* $q_H \geq q_M \geq q_L$*.*

Proof. Point 1 follow from the solution characterized in Propositions 6 and 7 (for convenience the quantities are reported in Table 1). The first part of Point 2 ($IC_{HM}(M)$ always binds) follows from Lemma 4 (when $\lambda = 0$) and Lemma A11 (when $\lambda > 0$). The second part follows from the fact that $IC_{ML}(M)$ can be assumed to hold as an equality. Suppose $\omega_{ML}(M) > \Delta\theta q_L(M)$. Then can decrease $\omega_{ML}(M)$ so that this holds as an equality. No constraint is violated and the profit of the monopolist is unaffected. Similarly, we show that $IC_{HM}(L)$ can be assumed to hold as an equality, implying $\omega_{HM}(L) = \Delta\theta q_M(L)$. Point 3 follows from incentive compatibility constraints for the second (terminal) period. We now turn to Point 4. From the fact that in the solution to the *WR-problem*, $q_H = \theta_H$ and the fact that (as shown in Propositions 6 and 7) $q_i \leq \theta_i$ for $i = H, M, L$, we have $q_H \geq q_i$ $i = M, L$. We, therefore, only need to prove that $q_M \geq q_L$. We will show this result case by case for all regions A1, A2, B1, B2 and B3. In cases $A1$ and $A2$, from (3.39) we have $q_M \geq q_L$ if and only if

$$1 - \frac{\mu_H}{\mu_M} + \frac{\mu_H + \mu_M}{\mu_L} \geq 0,$$

that is, $\frac{1}{\mu_L} \geq \frac{\mu_H}{\mu_M}$. In regions A1 and A2 we have $\mu_M \geq \mu^*(\alpha)$, as defined in Lemma 4. This condition can be written as

$$\frac{1}{\mu_L} \geq \frac{\mu_H}{\mu_M} + \delta\frac{3\alpha-1}{2}\frac{\mu_H}{\mu_M}\frac{3\alpha-1}{2\alpha} \quad \text{and} \quad \frac{1}{\mu_L} \geq \frac{\mu_H}{\mu_M} + \delta\frac{3\alpha-1}{2}\left(\frac{1-\alpha}{1+\alpha} + \frac{\mu_H}{\mu_M}\frac{3\alpha-1}{1+\alpha}\right).$$

clearly implying $\frac{1}{\mu_L} \geq \frac{\mu_H}{\mu_M}$. For case $B3$, we show in Proposition 7 that $q_M = q_L$. We now show that in regions $B1$ and $B2$ we have $q_M \geq q_L$ as well. In these region we have $\mu \in [\mu^{**}(\alpha), \mu^{*}(\alpha)]$. We have $q_M \geq q_L$ if and only if $1 - \frac{\mu_H - \lambda}{\mu_M} + \frac{\mu_H + \mu_M + \lambda}{\mu_L} \geq 0$. First order conditions in Lemma 4 clearly show that $\lambda > 0$ implies $\lambda_{HM}(L) > 0$, thus, $\omega_{HM}(L) = \Delta\theta q_M(L)$. Therefore, we have in regions B1 and B2,

$$q_M + \delta\frac{3\alpha-1}{2}q_M(M) = q_L + \delta\frac{3\alpha-1}{2}q_M(L).$$

When $\mu_M \geq \mu_0(\alpha)$, substituting optimal values (summarized in Table 1) we have

$$1 - \frac{\mu_H - \lambda_1}{\mu_M} + \frac{\mu_H + \mu_M + \lambda_1}{\mu_L} + \delta\frac{3\alpha-1}{2}\left[\frac{\lambda_1}{\mu_L}\frac{3\alpha-1}{1-\alpha} - \frac{\mu_H - \lambda_1}{\mu_M}\frac{3\alpha-1}{2\alpha}\right] = 0.$$

That can be re written as:

$$\left(1 - \frac{\mu_H - \lambda_1}{\mu_M} + \frac{\mu_H + \mu_M + \lambda_1}{\mu_L}\right)\left(1 + \delta\frac{(3\alpha-1)^2}{4\alpha}\right) =$$

$$\delta\frac{(3\alpha-1)^2}{4\alpha}\left[1 + \frac{\mu_H + \mu_M}{\mu_L} - \frac{\lambda_1}{\mu_L}\frac{3\alpha-1}{1-\alpha}\right].$$

We know from (3.57) that right hand side of the above equation is non-negative. Thus, $1 - \frac{\mu_H - \lambda_1}{\mu_M} + \frac{\mu_H + \mu_M + \lambda_1}{\mu_L} \geq 0$.

When $\mu_M < \mu_0(\alpha)$, substituting optimal values again (see Table 1) we have

$$1 - \frac{\mu_H - \lambda_2}{\mu_M} + \frac{\mu_H + \mu_M + \lambda_2}{\mu_L} + \delta\frac{3\alpha-1}{2}\left[\frac{\lambda_2}{\mu_L}\frac{3\alpha-1}{1-\alpha} - \frac{\mu_H - \lambda_2}{\mu_M}\frac{3\alpha-1}{1+\alpha} - \frac{1-\alpha}{1+\alpha}\right] = 0.$$

That can be rewritten as:

$$\left(1-\frac{\mu_H-\lambda_2}{\mu_M}+\frac{\mu_H+\mu_M+\lambda_2}{\mu_L}\right)\left(1+\delta\frac{(3\alpha-1)^2}{2(1+\alpha)}\right)$$

$$=\delta\frac{\alpha(3\alpha-1)}{1+\alpha}\left[\begin{array}{c}1+\frac{\mu_H+\mu_M}{\mu_L}\frac{3\alpha-1}{2\alpha}\\-\frac{\lambda_2}{\mu_L}\frac{3\alpha-1}{1-\alpha}\end{array}\right].$$

We know that (3.57) is always verified in the relevant range. Using this condition we can see that right hand side of the above equation is non-negative. Thus, we we have $1-\frac{\mu_H-\lambda_2}{\mu_M}+\frac{\mu_H+\mu_M+\lambda_2}{\mu_L}\geq 0$. ■

Consider the first period constraints. To show that IC_{LM} holds it is sufficient to prove:

$$\begin{aligned}0 &= U_L \geq \theta_L q_M+\delta\left[\alpha u_L(L)+\frac{1-\alpha}{2}u_M(M)+\frac{1-\alpha}{2}u_H(M)\right] \\ &= U_M-\Delta\theta q_M-\delta\frac{3\alpha-1}{2}u_L(M) \\ &= U_M-\Delta\theta q_M-\delta\frac{3\alpha-1}{2}q_L(M)\end{aligned} \tag{3.67}$$

Since $U_M=\Delta\theta q_L+\delta\frac{3\alpha-1}{2}q_L(L)$, (3.67) can be written as:

$$q_M+\delta\frac{3\alpha-1}{2}q_L(M)\geq q_L+\delta\frac{3\alpha-1}{2}q_L(L)$$

The fact that this inequality is satisfied follows from Point 1 and 4 in Lemma A14. (In the following, when we mention a point, we refer to the points of Lemma A14.)

Next, we show that IC_{MH} holds. From IC_{HM} we have:

$$U_H=U_M+\Delta\theta q_M+\delta\frac{3\alpha-1}{2}\left[u_H(M)-u_M(M)\right]$$

Thus,

$$\begin{aligned}
U_M &= U_H - \Delta\theta q_M - \delta\frac{3\alpha-1}{2}\left[u_H(M) - u_M(M)\right] \\
&= U_H - \Delta\theta q_H - \delta\frac{3\alpha-1}{2}\left[u_H(H) - u_M(H)\right] \\
&+ \Delta\theta(q_H - q_M) + \delta\frac{3\alpha-1}{2}\left[(u_H(H) - u_M(H)) - (u_H(M) - u_M(M))\right] \\
&> U_H - \Delta\theta q_H - \delta\frac{3\alpha-1}{2}\left[u_H(H) - u_M(H)\right].
\end{aligned}$$

The last inequality follows from the observation that:

$$u_H(H) - u_M(H) \geq \Delta\theta q_M(H) = \Delta\theta\theta_M > \Delta\theta q_M(M) = u_H(M) - u_M(M), \quad (3.68)$$

where the first inequality follows from the definition of $u_i(H)$, the first equality and the second inequality follow from Point 1. From (3.68) and the fact that $q_H > q_M$ (Point 4), it follows that IC_{MH} holds. We now turn to IC_{LH}. Using IC_{LM} first and then IC_{MH}, we have:

$$\begin{aligned}
U_L &\geq U_M - \Delta\theta q_M - \delta\frac{3\alpha-1}{2}\left[u_M(M) - u_L(M)\right] \\
&\geq U_H - \Delta\theta q_H - \delta\frac{3\alpha-1}{2}\left[u_H(H) - u_M(H)\right] - \Delta\theta q_M - \delta\frac{3\alpha-1}{2}\left[u_M(M) - u_L(M)\right] \\
&= U_H - 2\Delta\theta q_H - \delta\frac{3\alpha-1}{2}\left[u_H(H) - u_L(H)\right] \\
&+ \Delta\theta\left(q_H - q_M\right) + \delta\frac{3\alpha-1}{2}\left[(u_M(H) - u_L(H)) - (u_M(M) - u_L(M))\right] \\
&> U_H - 2\Delta\theta q_H - \delta\frac{3\alpha-1}{2}\left[u_H(H) - u_L(H)\right],
\end{aligned}$$

The last inequality follows from the observation that:

$$u_M(H) - u_L(H) \geq \Delta\theta q_L(H) = \Delta\theta\theta_L \geq \Delta\theta q_L(M) = u_M(M) - u_L(M), \quad (3.69)$$

where the first inequality follows from the definition of $u_i(H)$, the first equality and the second inequality follow from Point 1. From (3.69) and $q_H > q_M$ (Point 4), it follows that IC_{LH} holds.

Consider now the second period constraints. The constraints $IR_L(M)$, $IR_L(L)$ $IR_L(H)$, $IC_{ML}(H)$, and $IC_{HM}(H)$) follow immediately by the definition of the utilities at $t = 2$. The proof that $\langle \mathbf{U}, \mathbf{q} \rangle$ solves the seller's problem is therefore completed if we prove that it satisfies the constraints in the last two lines of (3.66). This result follows from the fact that the local downward incentive constraints are satisfied in period 2 and quantities are weakly monotonic after any history (Point 3). Finally, to see that the contract is optimal, we note that it maximizes expected profits in the less restricted *WR-problem*, so it must be optimal in the seller's problem. Note moreover that since the original problem is concave in q this is in fact the unique solution (in quantities). ■

3.16 From discrete to continuous types

In this section we show that the continuous case can be seen as the limit of the discrete case, so all problems of the FO-approach in the discrete version are inherited by the continuous version and viceversa. To keep the notation simple, we assume two periods and $u(\theta, q) = \theta q$. Consider a type set $\Theta = \left[\underline{\theta}, \overline{\theta}\right] \subset \mathbb{R}^{+}$, an associated prior distribution $\Gamma(\theta)$ at $t = 1$ and a conditional distribution $F(\theta' \,|\theta)$ at $t = 2$ defined on Θ. We assume $\Gamma(\theta)$ is differentiable in θ with density $\mu(\theta)$ and $F(\theta' \,|\theta)$ is differentiable in both θ, with derivative $F_\theta(\theta' \,|\theta)$, and θ', with density $f(\theta' \,|\theta)$. By standard methods we can obtain the following envelope formula (3.4):[45]

$$U'(\theta) = q(\theta) - \int_{\theta'} q(\theta' \,|\theta) \cdot F_\theta(\theta'|\theta)\, d\theta'$$

[45]See Baron and Besanko [1984], Besanko [1985], Laffont and Tirole [1996], Courty and Li [2000], Eso and Szentes [2007], and Pavan, Segal and Toikka [2013].

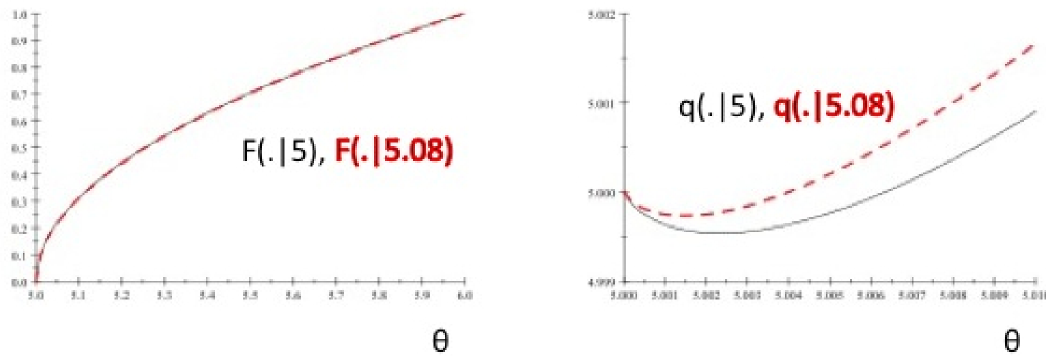

Figure 3.9: Example A1

and then derive the FO-optimal contract:

$$q(\theta' \,|\theta\,) = \theta' + \frac{1-\Gamma(\theta)}{\mu(\theta)} \frac{F_\theta\,(\theta'|\theta)}{f\,(\theta'|\theta)} \tag{3.70}$$

In the rest of this section, we refer to this as the *continuous model*. We start with an example.

Example A1. Assume $F(\theta' \,|\theta\,) = (\theta' - \underline{\theta})^{\gamma\theta}$ for $\theta' \in \left[\underline{\theta}, \overline{\theta}\right]$ where $\overline{\theta} - \underline{\theta} = 1$. If we assume that the prior on Θ is uniform, (3.70) implies that the FO-optimal contract is $q(\theta' \,|\theta\,) = \theta' + \frac{(\theta' - \underline{\theta})\ln(\theta' - \underline{\theta})}{\theta} \left(\overline{\theta} - \theta\right)$. Figure 3.9 plots the conditional distributions and the associated FO-optimal allocations when $\gamma = .1$ and $\Theta = [5, 6]$ after two histories $\theta_1 = 5.01$ and $\theta_1 = 5.08$. It is evident that the contract is non-monotonic in the realization at $t = 2$, θ_2. It is easy to see that this FO-optimal contract violates global contraints at $t = 2$ and so it is not incentive compatible.

We now explore the connection between the continuous model and the discrete model studied in the previous sections. The continuous model can be easily derived as the limit of the discrete model of the previous sections as follows. Define $\Theta^N = \{\theta_0, ..., \theta_N\}$ with $\theta_0 = \overline{\theta}$, $\theta_N = \underline{\theta}$ and $\theta_i = \theta_{i+1} + \Delta\theta_N$; and let $\Gamma^N\left(\theta_i\right) = \Gamma\left(\theta_i\right)$ and $F^N(\theta_j \,|\theta_i\,) = F(\theta_j \,|\theta_i\,)$. Given this, the probability of a type j at $t = 1$ is $\mu_j^N = \Gamma^N\left(\theta_j\right) - \Gamma^N\left(\theta_{j+1}\right)$ and the probability of a type i at $t = 2$ after a type j at

$t=1$ is $f^N(\theta_j|\theta_i) = F^N(\theta_j|\theta_i) - F^N(\theta_{j+1}|\theta_i)$.[46] In the rest of the section, we refer to this as the *discrete model*.

Consider a sequence of supports Θ^N for $N \to \infty$ such that $\Delta\theta_N \to 0$ as $N \to \infty$ and $\Theta^N \subseteq \Theta^{N+1}$, so that along the sequence the finite approximation of Θ becomes increasingly fine.[47] Using the formula derived in the paper (3.9), we can write the FO-optimal contract along the sequence as:

$$q_N(\theta_j|\theta_i) = \theta_j - \frac{\frac{1-\Gamma^N(\theta_i)}{\mu_i^N}}{} \frac{F^N(\theta_j|\theta_i) - F^N(\theta_j|\theta_{i-1})}{f^N(\theta_j|\theta_i)} \Delta\theta_N \tag{3.71}$$

for any $\theta_j \in \Theta^N$, $\theta_i \in \Theta^N$. Note that μ_i^N can be written as: $\mu_i^N = \frac{\Gamma(\theta_j)-\Gamma(\theta_{j+1})}{\Delta\theta_N} \cdot \Delta\theta_N$. and $f^N(\theta_j|\theta_i) = \frac{F^N(\theta_j|\theta_i)-F^N(\theta_{j+1}|\theta_i)}{\Delta\theta_N}\Delta\theta_N$. We can therefore rewrite (3.71) as:

$$q_N(\theta_j|\theta_i) = \theta_j + \left(1-\Gamma^N(\theta_i)\right) \frac{\left[F^N(\theta_j|\theta_i) - F^N(\theta_j|\theta_{i-1})\right]/\Delta\theta_N}{\left[\frac{\Gamma(\theta_i)-\Gamma(\theta_{i+1})}{\Delta\theta_N}\right]\left[\frac{F^N(\theta_j|\theta_i)-F^N(\theta_{j+1}|\theta_i)}{\Delta\theta_N}\right]}$$

This condition immediately implies that

$$\lim_{N\to\infty} q_N(\theta_j|\theta_i) = \theta_j + \frac{1-\Gamma(\theta_i)}{\mu(\theta_i)} \frac{F_\theta(\theta_j|\theta_i)}{f(\theta_j|\theta_i)} = q(\theta_j|\theta_i)$$

since $\mu_i^N/\Delta\theta_N \to \mu(\theta_i)$ and $f^N(\theta_j|\theta_i)/\Delta\theta_N \to f(\theta_j|\theta_i)$ as $N \to \infty$. It follows that the limit of the discrete FO-optimal contracts is equal to the continuous FO-optimal contract.[48]

This discussion makes it clear that there is a natural connection between discrete

[46] In both definitions, we are implicitly assuming a dummy "$N+1$" type with mass 0.

[47] For example, consider the sequence $(\theta_0^m, ..., \theta_N^m)$ such that $\theta_0^m = \underline{\theta}$, $\theta_N^m = \overline{\theta}$, $\theta_i^m - \theta_{i-1}^m = (\overline{\theta} - \underline{\theta})/2^m$ and so $N^m = 2^m$.

[48] Since $\Theta^N \subseteq \Theta^{N+1}$, if $\theta_j \in \Theta^N$, $\theta_i \in \Theta^N$, then $\theta_j \in \Theta^M$, $\theta_i \in \Theta^M$ for $M \geq N$, so $\lim_{N\to\infty} q_N^*(\theta_i|\theta_j)$ is well defined. To extend the contract for points on the real line that do not appear in the sequence of approximations we can consider, for example, the sequence of linear interpolations of the discrete contract. It is immediate to verify that this is a sequence of equicontinuous curves that converges to (3.70).

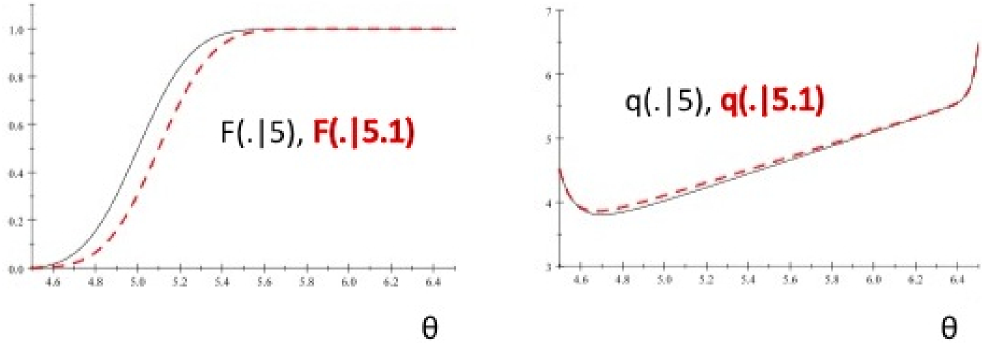

Figure 3.10: Example 8

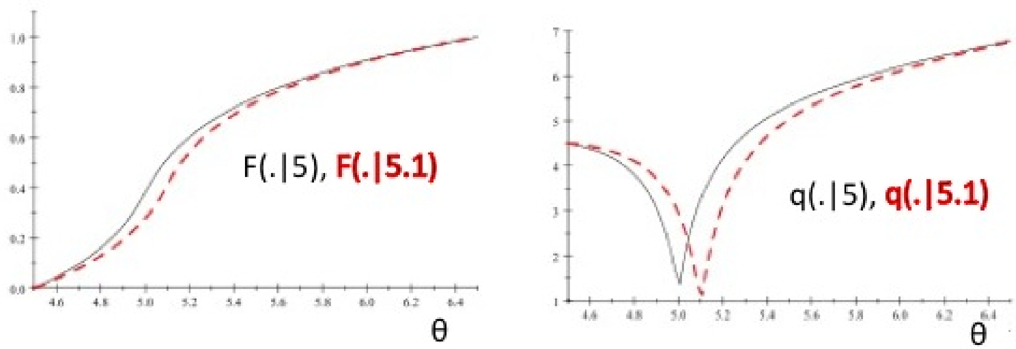

Figure 3.11: Example 9

and continuos types dynamic principal-agent models. In the light of this we can revisit the examples we have discussed in the previous sections in their continuos version.

Example 8 and 9 (cont.). Consider $f_\alpha(\theta'|\theta) = \alpha \cdot e^{-\frac{(\theta'-\theta)^2}{\sigma_\theta(\alpha)}}$ and $f_\alpha(\theta'|\theta) = \frac{\alpha}{1+\sigma_\theta(\alpha)|\theta'-\theta|}$ with all other parameters same as before. Note that $\sigma_\theta(\alpha)$ is chosen so that the probabilities sum to one. The larger is α, the higher is the persistence of the types. Figures 3.10 and 3.11 show two sample distributions and the associated quantities in period 2 that were plotted for the discretized case in Figures 5.3 and 5.4, respectively,. The contract is non-monotonic in two ways: first, for a given history, it is non-monotonic in θ_2. Because of this alone, the FO-optimal contract is not implementable and violates a global constraint. In addition to this, the FO-optimal contract is not monotonic with respect to θ_1; this can be seen from the fact that the contracts with the two different histories cross each other.

Bibliography

[1] Athey, S. and I. Segal (2013), "An Efficient Dynamic Mechanism," *Econometrica, forthcoming.*

[2] Baron, D. and D. Besanko (1984), "Regulation and Information in a Continuing Relationship," *Information Economics and Policy*, 1(3), 267-302.

[3] Battaglini, M. (2005), "Long-term Contracting with Markovian Consumers." *American Economic Review*, 95, 637–658.

[4] Battaglini, M. (2007), "Optimality and Renegotiation in Dynamic Contracting," *Games and Economic Behavior*, 60 (2), 213-246.

[5] Battaglini M. and S. Coate (2008), "Pareto Efficient Income Taxation with Stochastic Abilities," *Journal of Public Economics*, 92 (3-4), 844-868.

[6] Bergemann D. and J. Valimaki (2010). "The Dynamic Pivot Mechanism," *Econometrica*, 78(2), 771-789.

[7] Besanko, D. (1985), "Multiperiod Contracts between Prinipal and Agent with Adverse selection," *Economic Letters*, 17, 33-37.

[8] Biehl, A., (2001), "Durable-goods monopoly with stochastic values," *Rand Journal of Economics*, 32, 565–577

[9] Boleslavsky R. and M. Said (2013), "Progressive Screening: Long-Term Contracting with Privately Known Stochastic Process," *Review of Economic Studies*, forthcoming.

[10] Bolton, P. and M. Dewatripont (2005). *Contract Theory*. The MIT Press, Cambridge, MA.

[11] Chassang, S. (2013), "Calibrating Incentive Contracts," *Econometrica*, *forthcoming.*

[12] Courty, P. and H. Li (2000), "Sequential Screening," *Review of Economic Studies*, 67 (4), 697-718.

[13] Dewatripont, M. (1989), "Renegotiation and Information Revelation over Time: The Case of Optimal Labor Contracts." *Quarterly Journal of Economics*, 104 (3), 589–619.

[14] Eso, P. and B. Szentes (2007), "Optimal Information Disclosure in Auctions and the Handicap Auction," *Review of Economic Studies*, 74 (3), 705-731.

[15] Eso, P. and B. Szentes (2013), "Dynamic Contracting: An Irrelevance Result," *working paper.*

[16] Farhi, E. and I. Werning (2013), "Insurance and Taxation over the Life Cycle, " *Review of Economic Studies*, 80, 596-635.

[17] Garrett, D. and A. Pavan (2012), "Managerial Turnover in a Changing World," *working paper.*

[18] Golosov, M., M. Troshkin, and A. Tsyvinski, (2013), "Redistribution and Social Insurance," *working paper.*

[19] Guvenen, F., F. Karahan, S. Ozkan, and J. Song, (2013a), "What Do Data on Millions of U.S. Workers Say About Labor Income Risk?" *working paper.*

[20] Guvenen, F., S. Ozkan, and J. Song, (2013b), "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, forthcoming.

[21] Hartline, J. (2012), "Approximation in Mechanism Design," *American Economic Review P&P.*

[22] Hoffmann F. and R. Inderst (2011), "Presale Information," *Journal of Economic Theory,* 146 (6), 2333–2355.

[23] Inderst R. and M. Ottaviani (2012), "How (not) to pay for advice: A framework for consumer financial protection," *Journal of Financial Economics*, 105 (2), 393-411

[24] Kapicka, M. (2013), "Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach," *Review of Economic Studies,* 80 (3), 1027-1054.

[25] Krahmer, D. and R. Strausz (2013), "The Benefits of Sequential Screening, " *working paper.*

[26] Laffont, J.J. and D. Martimort (2002), *The Theory of Incentives,* Princeton University Press, Princeton, NJ.

[27] Laffont J.J. and J. Tirole (1990): "Adverse Selection and Renegotiation in Procurement," *Review of Economic Studies*, 57 (4), 597–625.

[28] Laffont J.J. and J. Tirole (1996), "Pollution Permits and Compliance Strategies," *Journal of Public Economics*, 62 (1–2), 85–125.

[29] Madarasz, K. and A. Prat (2012), "Screening with an Approximate Type Space," *working paper.*

[30] Maestri, L. (2013), "Dynamic Contracting under Adverse Selection and Renegotiation, " *working paper.*

[31] Milgrom, P. (2004), *Putting Auction theory to Work,* Cambridge University Press.

[32] Mussa M. and S. Rosen (1978), “Monopoly and Product Quality,” *Journal of Economic Theory*, 18 (2), 301–317.

[33] Myerson R. (1981), “Optimal Auction Design, ” *Mathematics of Operations Research*, 6 (1), 58-73.

[34] Pavan, A., I. Segal, and J. Toikka (2014), “Dynamic Mechanism Design: A Myersonian Approach,” *Econometirca,* 82(2): 601-653.

[35] Rey, P. and Salanie, B. (1990), “Long-Term, Short-Term and Renegotiation: On the Value of Commitment in Contracting.” *Econometrica*, 58(3), pp. 597–619.

[36] Stole, L. (2001), “Lectures on the Theory of Contracts and Organizations," *mimeo*, The University of Chicago.

[37] Roberts, K. (1982) “Long-Term Contracts.” *unpublished paper.*

[38] Royden, H. (1988), *Real Analysis*, Prentice Hall, Third Edition.

[39] Rustichini, A. and A. Wolinsky.(1995), “Learning about Variable Demand in the Long Run.” *Journal of Economic Dynamics and Control*, 19(5–7), pp. 1283–92.

[40] Strulovici, B. (2011), “Contracts, Information Persistence, and Renegotiation," *mimeo.*

[41] Townsend, R. M. (1982), “Optimal Multiperiod Contracts and the Gain from Enduring Relationships under Private Information.” *Journal of Political Economy*, 90 (6), 1166–86.

[42] Williams, N. (2011), “Persistent Private Information," *Econometrica,* 79 (4), 1233–1275.

[43] Zhang, Y. (2009), "Dynamic Contracting with Persistent Shocks," *Journal of Economic Theory*, 144, 635-675.