UNIVERSIDAD DE COSTA RICA

INSTITUTO DE INVESTIGACIONES EN CIENCIAS ECONOMICAS

AN ESSAY ON OPTIMAL-STOPPING DYNAMIC MECHANISMS

ALLAN ROBERTO HERNANDEZ CHANTO

Ciudad Universitaria Rodrigo Facio, Costa Rica
2012

If people do not believe
that Mathematics is simple
it is only because they do not realize
how complicated life is
John Louis von Neumann.

Contents

Abstract

The present paper presents a dynamic mechanism design model under uniform truncated intervals of time -called market windows- and endogenous report of departure times. The paper extends the framework of previous works to explicitly introduce the trade-off faced by buyers when outside trading opportunities are available in parallel markets. The main results show that the incentive compatibility conditions satisfy the standard extension of the results found by Myerson to the dynamic framework; and that sellers follow an optimal stopping rule which takes into account the possibility sellers have to collect fines if buyers walk away from the institutionalized market before the claimed departure time. Moreover, it is shown that given the structure of the windows, seller never terminates the market unilaterraly at *inter windows* periods, and instead makes an allocation decision only within market windows.

Chapter 1

Introduction

A recurrent and important feature of settings in which collective decisions are made is that individuals' preferences are not publicly observable, and therefore that decision-makers should rely upon the revelation of such preferences by the same agents. How the information can be elicited, and how the revelation of such information impacts the final payoffs induced by that action, is known as the *mechanism design problem.* In few words, mechanism design theory tries to disentangle one of the biggest challenges in economics: how to arrange economic interactions so that, when everybody behaves selfishly, the result is something "acceptable" to all. The notion of *m*echanism refers then to the institutions and the rules of the "game" that governs economic activities, whose span can range from the regulation structure of a particular market, to a policy for designating votes in a political election.

Defined in that way, mechanism design is at the heart of economics discipline, and since the pioneering work of Hurwitz (1960), Myerson (1981) and Maskin (1977 published 1999), many developments in the theory have permitted its application to a myriad of problems in auction and contract theory, assignation resources, social choice, political economy, public finance and monetary economics.

After the success of the theory to deal with problems of private information in static, non-collaborative environments, theorists posed themselves the question over how to extend its results to dynamic frameworks. Albeit some results can be easily transferred from the static framework, others are not straightforward but have new challenges embedded, especially regarding the way information is revealed over time, and with respect to the dimension of agents' type space (with many inter-temporal considerations included). In particular, the latter aspect makes incentive compatibility more difficult to analyze because there are new directions at which agents can deviate.

The intrinsic complication of these new models, far for being a theoretical artifact, finds empirical justification in many real-life situations. Indeed, the advent of technology has permitted that many mechanisms could be operated online rather than offline. As an example, online auctions are now a common way to trade, since it permits a seller to meet many possible buyers (with very distinct valuations over the object) in a very short interval of time. As usual, the objective of the seller is to maximize the expected revenue, which depends on the arrival process of bidders, and their stochastic valuations over the object at hand. However, auctions are not the only application of dynamic mechanisms. Other interesting applications include the operation of tenders in public projects, the allocation of computer resources (CPU, bandwidth) to a stream of tasks, the operation of the international market of human organs, and the running of soccer players' transfer market every summer in Europe. Many of these applications share roughly the same theoretical framework. Namely, an allocation of a single-unit object should be made for a stream of offers that arrive online, which in turn depends on the realizations of several stochastic variables. Nevertheless, in some cases more than in others, the temporal considerations play an important role in the determination of the strategies. Thus, consider for example a buyer trying to get an airline ticket online. The seller might ignore the precise time the agent arrives to the mechanism to see the price posted, the buyer's maximum willingness to pay, and or the maximum date travelers are willing to wait. In the same line, suppose you have a company that needs to close a deal with a party in order to sign a contract with a third party (e.g. suppose a builder that has to buy insurance in order to meet the requirements to gain a public tender). In such situation, the company is constrained by a deadline that affects its bidding behavior in the first market. Similarly, consider a soccer team bidding for a player in the summer market, but which has to start its pre-season at a specific date, which sets a deadline up to which the team is willing to wait for the player.

In all these examples deadline constraints are exogenous to seller's decision, but might be the case that agents themselves want to report a deadline endogenously if there is an outside opportunity that affects the latest time they are willing to sustain an offer. In the same example of the soccer team, consider a team looking for a specific position player in the European market, but with scouts looking for in alternative -less developed- markets. The team could decide not to sustain its offer very long for the player "auctioned" in the summer market, if it expects to find a suitable deal in those less known "parallel" markets. Moreover, imagine that the club has to sign a pre-agreement contract (and pay a fee) to enter the bid for the player, which the club lost if decides to walk away it in case of winning. Such contract would act as a hidden cost, which could be higher as nearer the date of disruption is from the close of the market window in that period.

Precisely, the fact that in some markets the time at which the item can be allocated is truncated in uniform intervals, introduce new theoretical challenges over the optimal way the planner shall behave, since delay decisions carry higher costs. In some environments, the problem for the seller is when to stop the market to maximize revenue, given all these dynamic elements in the strategic behavior of agents.

The present paper aims to explore the design of a mechanism that takes into account explicitly the presence of outside options in parallel markets, which make buyers strategic to report their departure time. Furthermore, the model considers a truncated market in a series of windows, to analyze how this structure impacts the seller's maximization problem through an optimal stopping rule. The work is organized as follows. Chapter 2 reviews the relevant literature on dynamic mechanisms with special emphasis on implementation and incentive compatibility matters. Chapter 3 presents a model for the design of a mechanism under truncated market windows and optimal report of departure times, accompanied by the main results and findings. Finally, chapter 4 presents some possible extensions of the baseline model, and chapter 5 concludes.

Chapter 2

Static Mechanism Design

The origins of modern mechanism design theory goes back to 1930s, when Oscar Lange and Abba Lerner were confronted with Friedrich von Hayek and Ludwig von Mises in the realm of the so-called *Planning Controversy.* One on side, Lange and Lerner argued that if public policies were executed "rightly," then central planning could replicate the performance of the free market, and sometimes surpasses markets in terms of efficiency. On the other side, von Hayek and von Misses argued that market "success" was not attainable by any other centralized mechanism (Maskin, 2008). At the center of this fascinating discussion were many scholars, and among them Leonid Hurwicz, who first noticed that the lack of rigorousity in the discussion was a real hindrance to try to reach robust conclusions about what approach was better. So, he entered the picture with his two pioneering (1960, 1972) works that laid the theoretical foundations of the field. Specifically, Hurwicz defined without ambiguity the key-stone concepts of "decentralization," "goodness," and "policies," at the time he incorporated the necessary elements of game theory and mathematical programming to characterize optimal solutions.

According to Maskin (2008), the work conducted by Hurwicz and his colleagues has produced a broad consensus among economists that von Hayek and von Mises where, in fact, correct where, (i) there are large number of buyers and sellers, so that no single agent has significant market power; and (ii) there are no significant externalities, that is an agent's consumption, production and information does not affect other's production and consumption. Nonetheless, Lange and Lerner were right if either assumption is violated.

Since this *momentum*, literature has bourgeoned following in general two different branches. On one hand, we have endeavors devoted to use structure settings to study particular environments and questions such as: how to allocate public goods, how to design auctions, how to conduct contests, how to design political elections, and how to structure contracts. On the other hand, there is a research line concerned primarily in proving results under fairly general conditions. Specifically, researchers here work in relaxing as much as possible the conditions about preferences, technologies and resources, in order to make main results more robust.

The second branch is exclusive of micro theorists, while the applied one spans many fields in economics such as contract theory, auction designs, international policy and political elections. Much of the relevant issues when analyzing different environments, come from the implementation side of theory. Namely, given a social goal, the task is to find a *mechanism* (or a set of institutions and rules) whose *predicted outcomes* (i.e. the set of equilibrium outcomes obtained) coincide with the *desirable* outcomes, according to the predetermined goal. Nonetheless, the nature of an "outcome" certainly depends on the problem at hand. Thus, for a state (county) seeking to elect a governor (mayor), an outcome is simply the choice of a candidate for that slot. Likewise, for a seller trying to assign an object through an auction, the outcome corresponds to an allocation specification of the object across potential buyers, along with the payments those buyers have to transfer to the seller. Another environment in which mechanism design appears naturally is in the provision of public goods. Here, an outcome comprises the quantities provided of such goods (public education, highways, recreational parks, among others), together with the arrangements by which they are financed.

In the same line of reasoning, the criteria of "optimality" or even "desirability" depends on the environment as well. For instance, in evaluating public goods the standard rule is the maximization of the *net social surplus*. Meanwhile, for electing governors, the property that a candidate would defeat every other competitor in a direct confrontation (i.e a *Condorcet* winner) is normally considered a dictum. In the auction scheme there are two criteria by which an outcome is typically analyzed: (i) whether the object is given to the agent who value it most (i.e. if the allocation is *efficient*); and (ii) whether the seller raises the maximum revenue possible. Notice that if resale is permitted and the object is not allocated to the most-caring potential buyer, then this agent and the current owner can be better off by trading between them (i.e. they can find an improving-pareto allocation).

Since a mechanism can be viewed as a game or as a set of rules and institutions, clearly it does matter who chooses the mechanism (i.e. who is the mechanism designer), which depends again on the setting of the problem. In the case of public goods, the designer is normally the government that chooses the structure through which the selected level of public good is financed. Similarly, in an auction the seller normally chooses the format on which it is run, determining specific rules for declaring winners and collecting payments. In the case of national political elections, by contrast, a mechanism is an *electoral procedure* prescribed regularly by country's constitution.

Notice that, if the planner (or mechanism designer) had perfect knowledge over individual preferences, then it would be able to "mandate" the optimal solution. Thus, in an auction setting the seller could simply award directly the object to buyer who value the object most, provided the buyer pays the amount established beforehand (which can be zero). Analogously, if a government has to invest in a public project (e.g. a bridge), then under perfect knowledge of preferences, it will be built if the total surplus of the population surpasses its cost. So, the difficulty in the problem -which is the main reason why mechanism design theory exists- is that normally the planner does not know agents' preferences, and these have to be elicited before making a decision. The problem is exacerbated by the fact that individuals (who of course know their own preferences and resources) have their own objectives, and so may not have the incentive to behave in a way that reveals what they know. Hence, one of the key properties of a mechanism must be *incentive compatibility*, that is the capability of the planner to induce individuals to reveal their private information accurately, through the design of particular institutions.

Following Maskin (2008), mechanism design theory is centered in answering three basic questions:

i) When is it possible to design incentive-compatible mechanisms for attaining social goals?

ii) What form might these mechanisms take when they exist?

iii) When is it possible to rule out the existence of such mechanisms on a theoretical basis?

Much of the literature developed in the last 30 years has attempted to provide detailed answers to these questions.

2.1 A Journey on the Development of MD theory

In his celebrated paper of 1960, Hurwicz defined a mechanism as a communication system in which participants send messages to each other and perhaps to a *message center*, and a prespecified rule assigns an outcome, for every collection of received messages. Almost at the same time, William Vicrey wrote a classic paper in 1961, which introduced the famous Vickrey auction (second price auction). Even today (mainly in the realm of dynamic environments), Vicrey auction continues to enjoy a special place in the annals of mechanism design. Later, another masterpiece was added to the theory by John Harsanyi (1967, 1968a, 1968b) who developed the theory of games with incomplete information, in particular Bayesian games, through a series of seminal papers. Even not thought in terms of the MD theory exclusively, Harsanyi's work has proved to be at the core of mechanism design and its multiple extensions[1]. Another influential paper in this journey was Hurwicz' 1972, where he formally introduced the concept of *incentive compatibility*, which allowed to incorporate incentives in the strategy of players, and then opened up mechanism design to a myriad of applications. On their part, Clarke (1971) and Grooves (1973), extended Vickrey mechanism and helped to disentangle how incentive compatibility works in quasi-linear environments.

Nonetheless, probably the two major advances in the 1970s correspond to the development of the *Revelation Principle* and the *Implementation Theory.* The former outcome, developed first by Gibbard (1973), and later by Maskin and Myerson, allowed theorists to center the attention to direct mechanism, leaving the development of real-world mechanisms (which are normally indirect[2]) to mechanism designers and practitioners. With respect to implementation theory, the main concern was focused on the feasibility of designing mechanisms such that all its equilibria were optimal. Maskin (1999) gave the first general solution to this problem[3]. In 1980's and 1990's the theory made numerous advances, especially with respect to the applications to specific environments. In the last part of 1990's and in the new millennium, the majority of endeavors have been concentrated in generalizing the well-rooted results in the theory, and extending the main results to dynamic settings (which will be the subject of the next chapter).

2.2 Background Theory

As mentioned before, in many circumstances private individual preferences have to be aggregated into social preferences, and ultimately into a collective decision trough a process of information disclosure conducted by the same agents. The way in which such information can be elicited, and the extent to which the information revelation problem constrains the ways in which social decision can respond to individual preferences, is the topic of mechanism design.

Following, Mas Colell et al. (1995), consider a set of I agents who must make a collective choice from some compact set X of possible alternatives. However, prior to the choice, an agent i

[1]Actually, one of the implementation criterion is based on a Bayesian-Nash equilibrium of the correspondent game

[2]Think for example in an auction, a public tender or a fiscal policy

[3]Even though the paper was published until 1999, it circulated as working paper in 1970's.

privately observes his preferences over the alternatives in X. Formally, this is modeled through the assumption that agents receive a signal θ_i which determines their preferences (θ_i is often called the agent's type). The compact set of all possible types for agent i is denoted by Θ_i. Each agent i is assumed to be an expected utility maximizer, whose Bernoulli-utility function is given by $u_i(x, \theta_i)$, where as usual $u(\cdot)$ is non-decreasing, continuous, and strictly concave.
Since θ_i is only observed by agent i, the environment is one of incomplete information. Nonetheless, agent's types are drawn from a publicly known prior distribution F with correspondent density $f \equiv dF$. Moreover, $\{\theta_i\}_{i=1}^{I}$, and $\{u_i\}_{i=1}^{I}$ are also public knowledge.
Because agents preferences depend on the realization of their types $\theta = (\theta_1, \cdots \theta_I)$, it is natural to assume that a collective decision depends on the types as well. To capture this dependence formally, a *social choice function* is introduced.

Definition 1. *A social choice function is a function* $\phi : \prod_{i=1}^{I} \Theta_i \rightarrow X$, *that assigns a collective choice* $\phi(\theta_1, \cdots \theta_I) \in X$ *for any profile of types* $\theta \in \prod_{i=1}^{I} \Theta_i$.

In many environments, it is desirable that the social choice function satisfies the Paretian property, which is stated below.

Definition 2. *A social choice function* $\phi : \prod_{i=1}^{I} \Theta_i \rightarrow X$ *is Paretian if for no profile* $\theta = (\theta_1, \cdots, \theta_I)$ *there exists an* $x \in X$ *such that* $u_i(\phi(\theta), \theta) \leq u_i(x, \theta)$ *for all* i, *and* $u_i(\phi(\theta), \theta) < u_i(x, \theta)$ *for at least one* i.

If the social function $\phi(\cdot)$ is Paretian, any other alternative choice cannot make an agent better off without worsening the utility of at least another player in the game.
The problem faced by agents is that types θ_i's, are not publicly observable, and therefore for the social choice $\phi(\theta_1, \cdots \theta_I)$ to be chosen, when types are actually $\theta = (\theta_1, \cdots, \theta_I)$ each agent i must be relied upon to disclosure her type i. However, for a given function $\phi(\cdot)$, some agents would find optimal for them, not to reveal their types truthfully. The next three examples will make clear the incentive problem of true revelation for agents involved in the mechanism.

Example 1. Undertaking of a Public Project.
Consider a society in which I agents must decide whether to undertake a public project (for example, a bridge, a highway, or a dam), whose cost must be funded by the agents themselves. An outcome is a vector $x = (k, t_1, \cdots t_I)$ where $k = 1$ if the project is undertaken (k=0 otherwise), and $t_i \in \mathbb{R}$ is a monetary transfer from (to if $t_i > 0$) agent i. The cost of the project is c, and

so the set of feasible alternatives for the I agents is

$$X = \{(k, t_1, \cdots t_I) : k \in \{0, 1\}, t_i \in \mathbb{R} \ and \sum_i t_i \leq -ck\}.$$

The constraint $\sum_i t_i \leq -ck$ reflects the fact that there is no source of outside funding (i.e. $c + \sum_i t_i \leq 0$ if $k = 1$ and $\sum_i t_i \leq 0$ otherwise). We assume that agent of type θ_i has the following quasilinear Bernoulli utility function:

$$u_i(x, \theta_i) = \theta_i k + (\bar{m}_i + t_i)$$

where $\bar{m}_i$ is the initial endowment of the numeraire ("money") and $\theta_i \in \mathbb{R}$. Here, θ_i can be thought as agent i's willingness to pay for the public project. In this context, a social choice function $f(\theta) = (k(\theta), t_1(\theta), \cdots t_I(\theta))$ is ex-post efficient if

$$k(\theta) = \begin{cases} 1 & if \sum_i \theta_i \geq c \\ 0 & otherwise \end{cases}$$

and

$$\sum_i t_i(\theta) = -ck(\theta)$$

Suppose that agents wish to implement a social choice function that sastisfies the latter conditions, and in which an egalitarian contribution is followed. That is, one in which $t_i(\theta) = -(c/I)k(\theta)$.

To consider a simple example, let $\Theta_i = \{\bar{\theta}_i\}$ for $i \neq 1$, and let $\Theta_1 = [0, \infty)$. Suppose that $c > \sum_{i \neq 1} \bar{\theta}_i > c(I-1)/I$. In that sense, agent 1 is a pivotal because if $\theta_1 \leq c - \sum_{i \neq j} \bar{\theta}_i$ the project is not undertaken, but if $\theta_1 \geq c - \sum_{i \neq j} \bar{\theta}_i$ and $\sum_{i \neq j} \bar{\theta}_i - c(I-1)/I > 0$, then it is. Let us examine agent 1's incentive for truthfully revealing his type, when $\theta_1 = c - \sum_{i \neq 1} \bar{\theta}_i + \epsilon$, for $\epsilon > 0$. If agent 1 reveals his true preferences, the project will be built because

$$\left(c - \sum_{i \neq 1} \bar{\theta}_i + \epsilon \right) + \sum_{i \neq 1} \hat{\theta}_i > c$$

In such scenario, agent 1's utility would be

$$\begin{aligned} & \theta_1 + \bar{m}_1 - \tfrac{c}{I} \\ = \quad & \left(c - \textstyle\sum_{i \neq 1} \bar{\theta}_i + \epsilon\right) + \bar{m}_1 - \tfrac{c}{I} \\ = \quad & \left(\tfrac{c(I-1)}{I} - \textstyle\sum_{i \neq 1} \bar{\theta}_i + \epsilon\right) + \bar{m}_1 \end{aligned}$$

For ϵ small enough, $\left(\frac{c(I-1)}{I} - \sum_{i \neq 1} \bar{\theta}_i + \epsilon\right) + \bar{m}_1$ is less than m_1, which is the utility obtained if agent 1 claims instead $\theta_1 = 0$. As a result the project is not built. Notice that when agent 1 causes the project to be undertaken, he has a positive externality on the other agents. However, given that he fails to internalize this effect, he has an incentive to understate his benefit from the project.

Example 2. Allocation of a Single Unit of an Indivisible Private Good.
Consider a setting in which there is a single unit of an indivisible private object that should be allocated to one of I agents. Money transfers can also be made. An outcome is a vector $x = (y_1, \cdots, y_I, t_1, \cdots t_I)$ where $y_i = 1$ if the agent i gets the object ($y_i = 0$ otherwise), and t_i is a monetary transfer received by agent i. The set of feasible alternatives is then,

$$\begin{aligned} \mathbf{X} = \{(y_1, \cdots, y_I, t_1, \cdots t_I) : y_i \in \{0,1\} \textit{ and } t_i \in \mathbb{R} \textit{ for all } i, \\ \sum_i y_i = 1, \textit{ and } \sum_i t_i \leq 0\} \end{aligned}$$

We suppose that type θ_i's Bernoulli utilty function takes the quasilinear form

$$u_i(x, \theta_i) = \theta_i k + (\bar{m}_i + t_i)$$

where $\bar{m}_i$ is agent i's initial endowment of the numeraire (money) and $\theta_i \in \mathbb{R}$ is the agent i's valuation of the good. In this situation, a social choice function $f(\theta) = (y_1(\theta), \cdots, y_I(\theta), t_1(\theta), \cdots t_I(\theta))$ is ex post efficient if it always allocates the good to the agent who values it most, and if it involves no "waste" of money. That is, if for all $\theta \in \prod_i \Theta_i$,

$$y_i(\theta) = (\theta_i - \max\{\theta_1, \cdots, \theta_I\}) = 0 \textit{ for all } i$$

and

$$\sum_i t_i(\theta) = 0$$

.

Two special cases that derive from this general setting are bargaining schemes, and auction setting (the case of first price and Vickrey auctions will be covered in detail in a further section). Notice that in the two examples examined here, the planner (either the government that intends to undertake the public project, or the seller that tries to allocate the object) can ask the type directly to the potential buyers, in order to choose the outcome ruled by the social choice function. Nonetheless, in general it is necessary to consider other structures to implement a social choice, rather than directly asking the players to reveal their type. This could be achieved through the design of institutions which agents interact with, and which implicitly determine their behavior. The formal notion of such institutions is known as a mechanism.

Definition 3. *The mechanism $\Xi(S_1, \cdots, S_I)$ is a collection of I strategy sets $(S_1, \cdots S_I)$ and an outcome function $\psi : \prod_{i=1}^{I} S_i \to X$.*

In a nutshell, a mechanism is a rule that governs the way in which agents' actions (circumscribed to S_i for each i) turn into a collective decision.
Formally, the mechanism Ξ combined with the space of possible types $\prod_{i=1}^{I} \Theta_i$, the probability density f, and the family of Bernoulli-utility functions $\{u_i\}_{i=1}^{I}$, conform a game of incomplete information. A strategy for agent i in the game of incomplete information induced by Ξ, is a function $s_i : \Theta_i \to S_i$. That is, for any type, the agent selects an action in her correspondent space.
Loosely speaking, a mechanism Ξ implements the social choice function ϕ, if the Bayes-Nash equilibrium of the game induced by Ξ, yields the same outcome as ϕ, for each $\theta \in \prod_{i=1}^{I} \Theta_i$. The next definition states this formally.

Definition 4. *The mechanism $\Xi(S_1, \cdots, S_I)$ implements the social choice function $\phi(\cdot)$ if there is an equilibrium strategy profile $(s_1^*(\cdot), \cdots s_I^*(\cdot))$ of the game induced by Ξ such that $\psi(s_1^*(\cdot), \cdots s_I^*(\cdot)) = \phi(\theta_1, \cdots \theta_I)$ for all $(\theta_1, \cdots \theta_I) \in \prod_{i=1}^{I} \Theta_i$.*

It is necessary to call attention on two aspects of the latter definition. The first is that the equilibrium solution is not explicitly stated, and therefore, any consistent solution for a Bayesian game fits in such definition. Moreover, the concept is silent on which equilibrium to choose if the game has multiple equilibria. The only requisite is that the equilibrium played (whatever be its definition) be coherent with the outcome displayed by ϕ.

2.2.1 Direct Mechanisms

Notice that the implementable space of functions is very large, because it involves the consideration of any kind of mechanism. Fortunately, a result derived by Myerson (1986), known as the *revelation principle*, shows that, to accomplish this daunting task, it is only necessary to take into account direct mechanisms: those in which each agent is asked to report her type, and given the announcements $(\hat{\theta}_1, \cdots \hat{\theta}_I)$, the alternative $\phi(\hat{\theta}_1, \cdots \hat{\theta}_I) \in X$ is chosen.

Definition 5. *A direct revelation mechanism, is a mechanism in which $S_i = \Theta_i$ for all i, and in which $\psi(\theta) = \phi(\theta)$ for all $\theta \in \prod_{i=1}^{I} \Theta_i$.*

Furthermore, the *revelation principle* shows that it is possible not only restrict the attention to direct mechanisms, but on those in which telling the truth is an optimal strategy for each agent. When that is possible, it is said that the mechanism is truthfully implementable.

Definition 6. *The social choice function, ϕ is truthfully implementable, or incentive compatible, if the direct revelation mechanism $\Xi(\Theta_1, \cdots \Theta_I, \psi(\cdot))$ has an equilibrium in which $s_i^*(\theta_i) = \theta_i$ for all $\theta_i \in \Theta_i$, and all i.*

Basically that definition says that ϕ is incentive compatible if telling the truth is the Bayes-Nash equilibrium of the game induced by Ξ.

2.2.2 Dominant Strategies Implementation

In the present incomplete information environment, strategy $s_i : \Theta_i \to S_i$ is a weakly dominant strategy for agent i in the mechanism $\Xi(S_1, \cdots S_I, \psi(\cdot))$, if for all $\theta_i \in \Theta_i$ and all possible strategies $s_{-i} \equiv (s_1, \cdots, s_{i-1}, s_{i+1} \cdots, s_I) \in \prod_{j \neq i} S_i \equiv S_{-i}$,

$$E_{\theta_{-i}}[u_i(\psi(s_i(\theta_i), s_{-i}(\theta_{-i})), \theta_i)|\theta_i] \geq E_{\theta_{-i}}[u_i(\psi(\hat{s}_i(\theta_i), s_{-i}(\theta_{-i})), \theta_i)|\theta_i] \tag{2.1}$$

for all $\hat{s}_i \in S_i$. Holding for all $s_{-i}(\cdot)$, and all θ_i, the latter is equivalent that for all $\theta_i \in \Theta_i$

$$u_i(\psi(s_i(\theta_i), s_{-i}), \theta_i) \geq u_i(\psi(\hat{s}_i, s_{-i}), \theta_i)$$

for all $\hat{s}_i \in S_i$ and all $s_{-i} \in S_{-i}$.

Definition 7. *The strategy profile $s^* = (s_1^*, \cdots s_I^*)$ is a dominant strategy equilibrium of mech-*

anism $\Xi(S_1, \cdots S_I, \psi(\cdot))$ *if for all i, all* $\theta_i \in \Theta_i$,

$$u_i(\psi(s_i^*(\theta_i), s_{-i}), \theta_i) \geq u_i(\psi(\hat{s}_i, s_{-i}), \theta_i)$$

for all $\hat{s}_{-1} \in S_i$, $s_{-i} \in S_{-i}$.

Focusing on the equilibrium concept, the next definition introduces the notion of dominant strategies equilibrium.

Definition 8. *The mechanism* $\Xi(S_1, \cdots S_I, \psi(\cdot))$ *implements* $\phi(\cdot)$ *in dominant strategies if there exists a dominant strategy equilibrium of* Ξ, $s^* = (s_1^*, \cdots s_I^*)$, *such that* $\psi(s^*(\theta)) = \phi(\theta)$ *for all* $\theta \in \Theta$.

Moreover, if the dominant equilibrium strategies is one of the truth-telling class, then $\phi(\cdot)$ is truthfully implementable in dominant strategies.

Definition 9. *The social choice function* $\phi(\cdot)$ *is truthfully implementable in dominant strategies if* $s_i^*(\theta_i)$=θ_i *for all* $\theta_i \in \Theta_i$, *all* $i = 1, \cdots, I$, *is a dominant strategy equilibrium of the direct mechanism* $\Xi(\Theta_1, \cdots \Theta_I, \phi(\cdot))$.

Given stated this, it is now possible to introduce the revelation principle in dominant strategies.

Proposition 1. *Suppose that there exists a mechanism* $\Xi(S_1, \cdots S_I, \psi(\cdot))$ *that implements the social function* $\phi(\cdot)$ *in dominant strategies. Then,* ϕ *is truthfully implementable in dominant strategies.*

Proof. If $\Xi(S_1, \cdots S_I, \psi(\cdot))$ implements ϕ in dominant strategies, then there exists a profile of strategies $s^* = (s_1^*, \cdots, s_I^*)$ such that $\psi(s^*(\theta)) = \phi(\theta)$ for all θ. Then, for all $\theta_i \in \Theta_i$.

$$u_i(\psi(s^*(\theta_i), \theta_{-i}), \theta_i) \geq u_i(\psi(\hat{s}_i, s{-}i), \theta_i) \tag{2.2}$$

for all $\hat{s}_i \in S_i$, $s_{-i} \in S_{-i}$. Condition (2.2) implies in particular that for all i, all $\theta_i \in \Theta_i$

$$u_i(\psi(s^*(\theta_i), s_{-i}^*(\theta_{-i}), \theta_i) \geq u_i(\psi(s^*(\hat{\theta}_i), s_{-i}^*(\theta_{-i}), \theta_i) \tag{2.3}$$

for all $\hat{\theta}_i \in \Theta_i$, all $\theta_{-i} \in \Theta_{-i}$. Since, $\psi(s^*(\theta)) = \phi(\theta)$ for all θ, (2.3) means that

$$u_i(\phi(\theta_i, \theta_{-i}), \theta_i) \geq u_i(\phi(\hat{\theta}_i, \theta_{-i}), \theta_i)$$

for all $\hat{\theta}_i \in \Theta_i$ $\theta_{-i} \in \Theta_{-i}$. That is, $\phi(\cdot)$ is truthfully implementable in dominant strategies. ■

Implementation in dominant strategies can be expressed alternatively in terms of alternatives' *lower countour sets.* Define a *lower countour set* of an alternative x when agents i has type θ_i, by

$$L_i(x, \theta_i) = \{z \in \mathbf{X} : u_i(x, \theta_i) \geq u_i(z, \theta_i)\}$$

Using, this notation, the set of social choice functions that can be truthfully implementable in dominant strategies is given by the following proposition.

Proposition 2. *The social choice function $\phi(\cdot)$ is truthfully implementable in dominant strategies if and only if for all i, all $\theta_{-i} \in \Theta_{-i}$, and all pairs of types for agent i, ${\theta'}_i$ and $\theta''_i \in \Theta_i$, we have*

$$\phi(\theta''_i, \theta_{-i}) \in L_i(\phi(\theta'_i, \theta_{-i}), \theta'_i)$$

and

$$\phi(\theta'_i, \theta_{-i}) \in L_i(\phi(\theta''_i, \theta_{-i}), \theta''_i)$$

Proof. See Mas Colell et al. (1995). ■

Another important result with respect to implementation in dominant strategies is the so-called Gibbard-Satterthwaite theorem, disovered independently by Gibbard (1973), and Satterthwaite (1975). It is a sort of impossibility result in the spirit of Arrow's theorem, that qualifies the class of mechanisms that can be implemented in dominant strategies. In general, it shows that for a very large class of problems, there is no hope of implementing satisfactory social choice functions in dominant strategies.

To delve into the necessary elements of the theorem, let $\mathcal{P}$ denote the set of all rational preference relations $\succeq$ on X having the property that not two alternatives are indifferent, and let $\mathcal{P}_i = \{\succeq : \succeq_i = \succeq_i (\theta_i) \text{ for some } \theta_i \in \Theta_i\}$ be agent i's set of possible ordinal preference relations over X. Moreover, we denote $\phi(\Theta) = \{x \in X : \phi(\theta) = x \text{ for some } \theta \in \Theta\}$. The following two properties of social choice functions, are also necessary to understand the main result of Gibbard and Satterthwaite.

Definition 10. *The social choice function $\phi(\cdot)$ is dictatorial if there is an agent i, such that, for all $\theta = (\theta_1, \cdots \theta_I) \in \Theta$*

$$\phi(\theta) \in \{x \in X : u_i(x, \theta_i) \geq u_i(y, \theta_i) \text{ for all } y \in X\}$$

In summary, a social choice function is dictatorial if it always chooses a top alternative for some agent i.

Definition 11. *The social choice function $\phi(\cdot)$ is monotonic if, for any θ, if θ' is such that $L_i(\phi(\theta), \theta_i) \subset L_i(\phi(\theta), \theta'_i)$, then, $\phi(\theta') = \phi(\theta)$*

Monotonicity requires the following: Suppose that $\phi(\theta) = x$, and that the I agents's types change to a $\theta' = (\theta'_1 \cdots \theta_I)$, with the property that no agent finds that some alternative that was weakly worse for him than x when his type was θ_i becomes strictly preferred to x when his type is θ'_i. Then x must still be the social choice. Once introduced such definitions, we now state the Gibbard-Satterthwaite theorem.

Theorem 1. *Suppose that X is finite and contains at least three elements, that $\mathcal{P}_i = \mathcal{P}$ for all i, and that $f(\Theta) = X$. Then the social choice function $f(\cdot)$ is truthfully implementable in dominant strategies if and only if it is dictatorial.*

Proof. See Mas Colell et al. (1995). ■

Given this negative conclusion, if we are to have any hope of implementing desirable social choice functions, we must either focus on more restricted environments, or weaken the demands of the implementation concept by allowing less robust equilibrium notions, such as Bayesian Nash equilibria, which will be reviewed in the next section.

2.2.3 Bayesian Implementation

Another important implementation concept is the Bayesian implementation, which relies on the knowledge of the joint distributions of types by all agents in the game. As it will be seen, it is a weaker solution than its dominant strategies counterpart. To discuss implementation issues, we begin by introducing the concept of a Bayesian-Nash equilibrium first.

Definition 12. *The strategy profile $s^* = (s_1^*, \cdots, s_I^*)$ is a Bayesian-Nash equilibrium of mechanism $\Xi(S_1, \cdots S_I, \psi(\cdot))$ if for all i, $\theta_i \in \Theta_i$,*

$$E_{\theta_{-i}}[u_i(\psi(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i)|\theta_i] \geq E_{\theta_{-i}}[u_i(\psi(\hat{s}_i(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i)|\theta_i] \tag{2.4}$$

for all $\hat{s}_i \in S_i$.

Definition 13. *The mechanism $\Xi(S_1, \cdots S_I, \psi(\cdot))$ implements the social function $\phi(\cdot)$ in Bayesian-Nash equilibrium, if there is a Bayesian-Nash equilibrium of Ξ, $s^* = (s_1^*, \cdots, s_I^*)$ such that $\psi(s^*(\theta)) = \phi(\theta)$ for all $\theta \in \Theta$.*

Analogously to the dominant strategies case, it will be seen that a social choice function ϕ is implementable in Bayesian equilibrium if it is truthfully implementable according to the following definition.

Definition 14. *The social choice function $\phi(\cdot)$ is truthfully implementable in Bayesian-Nash equilibrium, if $s_i^*(\theta_i)=\theta_i$ for all $\theta_i \in \Theta_i$, all $i = 1, \cdots, I$, is a Bayesian-Nash equilibrium of the direct mechanism $\Xi(\Theta_1, \cdots \Theta_I, \phi(\cdot))$.*

The ability to restrict the analysis (without loss of generality) to truthfully implementable social choice functions, is a consequence of the *revelation principle for Bayesian implementation.*

Proposition 3. *Suppose that there exists a mechanism $\Xi(S_1, \cdots S_I, \psi(\cdot))$ that implements the social function $\phi(\cdot)$ in Bayesian-Nash equilibrium. Then, ϕ is truthfully implementable in Bayesian-Nash equilibrium.*

Proof. If $\Xi(S_1, \cdots S_I, \psi(\cdot))$ implements ϕ in Bayesian-Nash equilibrium, then there exists a profile of strategies $s^* = (s_1^*, \cdots, s_I^*)$ such that $\psi(s^*(\theta)) = \phi(\theta)$ for all θ. Then, for all $\theta_i \in \Theta_i$.

$$E_{\theta_{-i}}[u_i(\psi(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i)|\theta_i] \geq E_{\theta_{-i}}[u_i(\psi(\hat{s}_i, s_{-i}^*(\theta_{-i})), \theta_i)|\theta_i] \tag{2.5}$$

for all $\hat{s}_i \in S_i$. Condition (2.5) implies in particular that for all i, all $\theta_i \in \Theta_i$

$$E_{\theta_{-i}}[u_i(\psi(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i)|\theta_i] \geq E_{\theta_{-i}}[u_i(\psi(s_i^*(\hat{\theta}_i, s^[*]_{-i}(\theta_{-i})), \theta_i)|\theta_i] \tag{2.6}$$

for all $\hat{\theta}_i \in \Theta_i$. Since, $\psi(s^*(\theta)) = \phi(\theta)$ for all θ, (2.6) means that

$$E_{\theta_{-i}}[u_i(\phi(\theta_i, \theta_{-i}), \theta_i)|\theta] \geq E_{\theta_{-i}}[u_i(\phi(\hat{\theta}_i, \theta_{-i}), \theta_i)|\theta_i]$$

for all $\hat{\theta}_i \in \Theta_i$. That is, $\phi(\cdot)$ is truthfully implementable in Bayesian equilibrium. ■

From the definitions, it is possible to notice immediately that the Bayesian implementation concept is a weaker notion than the dominant strategies mode. Since any dominant strategies equilibrium of the game induced by Ξ, is a Bayesian-Nash equilibrium of the same mechanism, any social function $\phi(\cdot)$ which is implementable in dominant strategies is also implementable in Bayesian-Nash equilibrium. The fact is that within Bayesian implementation, truth telling need only give an agent i her highest payoff averaging over all possible θ_i that might arise for the other agents, while the dominant strategy concept requires that truth telling be agent i's best strategy for every possible θ_{-i}.

A very popular setting in which mechanism design theory emerges naturally is in the administration of auctions, which can be seen as institutions built to replace over-the-counter markets. Example 2 showed how allocation of a single unit object can be addressed in the realm of theory discussed here. A natural option is to ask to potential buyers directly about their valuation over the object, and assign it to the agent who values it most. However, there are other indirect mechanisms that can be used to attain the same goal, but with the advantage of avoiding the perverse incentives agents have to hide the truth. The next section discuss in detail the main two auctions formats, namely, first price auction and the Vickrey auction.

2.3 First Price and Vickrey Auctions as Static Mechanisms

Modern auction theory is founded in the pioneer work of Vickrey (1961), who derived the main results of strategic behavior of buyers and sellers under different auction schemes, and introduced the celebrated equivalence revenue theorem, which states that under some general conditions, a seller will raise the same revenue, independently of the auction format utilized. Later advancements have introduced different assumptions over the information framework, among which the work of Myerson (1981), Riley and Samuelson (1981) and Milgrom (1989), stand out. For the sake to analyze the main properties of such mechanisms, models assume a single-object environment, and then extend some properties to multi-object auctions, in which combinatorial issues make the problem harder to be solved. Nonetheless, this section will be focused in the former class.

Suppose there is a single object for sale and I potential buyers who are bidding for the object at hand. A bidder i attaches a value of V_i to the object -which constitutes the maximum amount she is willing to pay for it. Each V_i is independently and identically distributed on some compact support $[V_l, V_u] \in \mathbb{R}_+$ according to a strictly increasing, right continuous, distribution F. For convenience, it is assumed that F admits a continuous density $f = dF$, and that -without loss of generality- $V_l = 0$. If the distribution happens to have the entire nonnegative real line as a full support, it is assumed that $\mathbb{E}[V_i] < \infty$.

Following the representation of a Bayesian game, bidder i knows the realization v_i of V_i, but only knows that other bidders' values are independently distributed with respect to F. To avoid risk considerations, bidders are considered risk-neutral agents whose aim is to maximize their expected value. That is, there is no bias to participate or not in a lottery *per se*, and only the final expected profit is taken into account. All information's structure, other than the realization of the valuations, is assumed to be public to all participants. Specifically, the number of bidders is considered fixed and known before hand by the seller, and the distribution function F is common knowledge.

At the time, the mechanism is considered free of transactional costs, in the sense agents do not have to pay a fee for participate in the auction, and the seller does not run any administrative cost to operate it. Moreover, liquidity and inter-temporal constraints are ruled out from the problem, thus, it is assumed that at any moment a bidder i has sufficient resources, up to her maximum valuation, to bid and pay -if she wins[4]- an amount of b_i. Since all bidders face the same distribution from which values are drawn, the environment is known as the symmetric bidders case. The implications of this theoretical framework will be examined on the two major simultaneous-move auctions format: first and second price sealed-bid auctions. In such schemes, all bidders are called to move simultaneously and share the same information set, unlike English and Dutch formats, where information sets are deviled sequentially, much in the spirit of a filtration. Albeit online auctions can be run sequentially, its underpinnings are those of the former.

2.3.1 Optimal Strategy for Bidders and Bayesian Nash Equilibrium (BNE)

As mentioned before, first and second price sealed-bid auctions conform a Bayesian game among bidders. Following Selten (1965), a strategy for a bidder is a function from the space of valuations to the nonnegative real line, namely $b_i : [0, V_u] \rightarrow \mathbb{R}_+$. Here, b_i is a decision rule that assigns to each type of bidder[5], an amount of money to compete for the object being auctioned. Since agents would follow symmetric strategies, the interest lies in finding out the properties of the symmetric equilibrium. Although first-price auction is a more intuitive format, the derivation of second price auction is presented before, because it is identifiable with an open ascending (English) auction, which is widespread in online mechanisms, and its equilibrium is easier to derive. Under that format, the payoff of bidder i who bids b_i can be summarized in

$$\pi_i = \begin{cases} v_i - \max_{j \neq i} b_j & if \quad b_i > \max_{j \neq i} b_j \\ 0 & if \quad b_i < \max_{j \neq i} b_j \end{cases} \tag{2.7}$$

[4] Strictly speaking, there are schemes known as all-pay contests in which all participants shall pay the bid, independently if they win the object. Nonetheless, here are considered cases in which only the winners pay. For more information on all-pay contests, see Siegel (2007).

[5] Determining a type by her maximum willingness to pay.

Like many models in literature, we assume that if there is a tie, the object is assigned to each winner bidder with equal probability. The following proposition shows the BNE bidder strategy in that format.

Proposition 4. *In a second price sealed-bid auction, $b_i(v_i) = v_i$ constitutes the unique BNE of the game.*

Proof. See Krishna (2002).

Once defined the optimal strategy a bidder would follow in this framework, it is possible to calculate the expected amount she would pay in equilibrium. As a matter of convenience, fix a bidder i and let the random variable $X_{(1),i}$ denote the maximum value among the $I-1$ remaining bidders (i.e. the highest statistic order of $\{V_j\}_{j\neq i}$). Let also H denote the distribution function of $X_{(1),i}$, which have h as its density. Given the assumption of independence, it is clear that for all x, $H(x) = F(x)^{I-1}$. According to this, the expected payment of a bidder in a second price format corresponds to

$$\begin{aligned} a_i^{II}(v_i) &= \text{Prob[Win]} \times \mathbb{E}[\text{2nd highest value} \mid v_i \text{ is the highest value}] \\ & H(v_i) \times \mathbb{E}[X_{(1),i} | X_{(1),i} < v_i] \end{aligned} \tag{2.8}$$

On the other hand, when the auction is run as a first-price, payoffs are given by

$$\pi_i = \begin{cases} v_i - b_i & if \quad b_i > \max_{j\neq i} b_j \\ 0 & if \quad b_i < \max_{j\neq i} b_j \end{cases}$$

As before, in the case of a tie, winners participate in a lottery in which all of them have equal probability to get the object.

Bidding their true value is not a weakly dominant strategy for buyers in a first price auctions, as it is in the second price scheme. Here, if the buyer wins the auction following this strategy, she will obtain a zero payoff, but if $\max_{j\neq i} b_j < v_i$, then any bid $b_i \in (\max_{j\neq i} b_j, v_i)$ would assure a strictly possitive payoff. So, given the decision rules of the rest of buyers, each individual faces a trade off in her own bidding strategy, because an increase in the bid will increase the probability of winning, but will decrease the gains from winning. Equilibrium strategies are summarized in the following proposition.

Proposition 5. *In a first price sealed bid auction $b_i(v_i) = \mathbb{E}[X_{(1),i} | X_{(1),i} < v_i]$ for all $i = 1, 2, \cdots I$ constitutes a BNE profile for this game.*

Proof. See Krishna (2002).

Unlike second-price auctions in which the winner pays the maximum bid of her opponents, in the first-price format bidders end up paying their own bids, if they win. Therefore, given the strategy stated in proposition 2, the expected payment of a bidder with value v is

$$a_i^I(v_i) = \text{Prob[Win]} \times \text{Amount of bid} = H(v_i) \times \mathbb{E}[X_{(1),i}|X_{(1),i} < v_i] \tag{2.9}$$

which is the same as in the second-price auction

Because the expected revenue of the seller is just the summation of the *ex ante* expected payments of the bidders, this is the same under both formats. This somewhat stunning result is a corollary of a general theorem known as the equivalence revenue theorem, firstly introduced by the harbinger work of Vickrey (1961), and extended later for Myerson (1981) to a broader class of mechanisms.

Indeed, the *ex ante* expected payment of a particular bidder in either format corresponds to

$$\mathbb{E}[a(V)] = \int_0^{V_u} a(v)f(v)dv = \int_0^{V_u} \left(\int_0^{v} xh(x)dx \right) f(v)dv \tag{2.10}$$

Applying Fubbini's theorem to the latter expression it is obtained that

$$\mathbb{E}[a(V)] = \int_0^{V_u} \left(\int_x^{V_u} f(v)dv \right) xh(x)dx = \int_0^{V_u} x(1 - F(x))h(x)dx \tag{2.11}$$

The revenue a seller expects to receive in such mechanism is I times the expected payment of an individual bidder, so it can be computed as:

$$\mathbb{E}[R] = I \times \mathbb{E}[a(V)] = N \int_0^{V_u} x(1 - F(x))h(x)dx \tag{2.12}$$

Nonetheless, even when sellers can expect the same revenue in both formats, they do not exhibit the same level of risk. In second price auctions, prices can vary between 0 and V_u, but in its first price counterpart they can only vary between 0 and $\mathbb{E}[X_{(1),i}]$. In that sense, the distribution of prices in the former case is a *mean preserving spread* of the latter. Therefore, in the case of risk-averse sellers, they would prefer the first-price format, under the assumption that buyers are risk-neutral (Krishna, 2006).

2.3.2 Reserve Prices

In the analysis conducted before, sellers do not control the administration of the auction, beyond to decide under which format it will be framed in. In particular, they just receive all offers and allocate the object according to each scheme's rule. However, since *a priori* an auction involves more uncertainty than a bilateral bargaining, or an over-the-counter mechanism, sellers might opt to impose a reserve price v^R under which the object is never sold.

Clearly, no buyer i whose valuation is lower than such reserve price, $v_i < v^R$, can make a positive profit in the auction, but for the rest of participants it has no impact on their equilibrium strategies. Thus, in a second price auction, bidding their true value is still a weakly dominant strategy for all bidders. Moreover, the expected payment of a bidder with value $v_i = v^R$ is $v^R H(v^R)$, and for those with $v_i > v^R$ is

$$a_i^{II}(v_i, v_R) = v^R H(v^R) + \int_{v^R}^{v_i} xh(x)dx \tag{2.13}$$

due to the winner will pay the reserve price v_R, whenever $\max_{j \neq i} b_j < v^R$. Analogously, in the first-price auction scheme, the equilibrium bidding strategy for any bidder i with value $v_i \geq v^R$ is given by

$$b_i(v_i) = \mathbb{E}[\max\{X_{(1),i}, v^R\} | X_{(1),i} < v_i] = v^R \frac{H(v^R)}{H(v_i)} + \frac{1}{H(v_i)} \int_{v^R}^{v_i} xh(x)dx \tag{2.14}$$

and her expected payment is given by the same expression as in (2.13).

Analogously to (2.10), the *ex ante* expected payment of a bidder is

$$\mathbb{E}[a(V, v^R)] = \int_0^{V_u} a(v, v^R) f(v) dv = v^R (1 - F(v^R)) H(v^R) + \int_{v^R}^{V_u} x(1 - F(x)) h(x) dx \tag{2.15}$$

As it was stated, the seller wants to maximize the expected benefits, and therefore will try to set the reserve price at the optimum level. Notice first that there is a tradeoff between the expected gains received, given that the object is effectively sold, and the probability of allocating the object to some bidder. As bigger the reserve price is, bigger is the probability that the maximum bid be lower than the threshold established. However, a big reservation price sets a higher minimum price at which the object could be sold, and therefore increases the potential revenue the seller would accrue if the object is actually sold to the maximum bidder. For the sake of pointing out the analysis of reserve prices in the actual context, let $v_0 \in [0, V_u)$ denote the value that the seller attaches to the object (i.e. the value of use that the object would render to the seller). Clearly, the seller will not set a reserve price v^R lower than v_0, since it would obtain less benefits than those she would obtain by keeping it. The overall payoff the seller expects by setting a reservation price $v^R \geq v_0$ is

$$\Pi_{v_0} = I \times \mathbb{E}[a(V, v^R)] + F(v^R)^I v_0$$

Taking the first derivative of Π_0 with respect to v^R, it is obtained that

$$\frac{d\Pi_{v_0}}{dv^R} = I[1 - F(v^R) - v^R f(v^R)]H(v^R) + IH(v^R)f(v^R)v_0$$

Recall that a hazard rate of a distribution F is defined as $\lambda(v) = f(v)/1 - F(v)$. Therefore, the latter expression can be rewritten as

$$\frac{d\Pi_{v_0}}{dv^R} = I[1 - (v^R - v_0)\lambda(v^R)(1 - F(v^R))H(v^R)$$

Notice that if $v_0 > 0$, the derivative of Π_0 at $v^R = v_0$ is $N(1 - F(v^R))H(v^R) > 0$, which implies that there is a room to increase the reserve price. On the other hand, if $v_0 = 0$, then by setting $v^R = 0$, the whole derivative becomes zero, but as long as the hazard rate is bounded, the expected payoff function reaches a local minimum at $v^R = 0$. As a consequence, a seller should always set a reserve price that exceeds v_0.
Setting the first order conditions, it turns out that the reserve price should satisfy:

$$\begin{aligned} v^{R^*} - v_0\lambda(v^{R^*}) &= 1 \Rightarrow \\ v^{R^*} - \frac{1}{\lambda(v^{R^*})} &= v_0 \end{aligned}$$

If the hazard rate $\lambda(\cdot)$ is increasing, the necessary condition is also sufficient, and it is independent of the number of participants. Therefore, the expected gain from setting a reserve price above the use of value for the seller (v_0) surpasses the expected loss, a result commonly known as the *exclusion principle*. It reflects the fact that it is optimal for the seller to exclude some bidders whose valuations lie in the range (v_0, v^R).
Even though the use of a reserve price could increase the revenue of the seller in any of the formats explored, it could have harmful effects on efficiency. To see this, consider an auction (second or first-price) without reservation prices, in which the seller attaches a zero value to

the object for sale. Clearly, the object will be always allocated to the agent who values it more (although the final price paid depends on the format itself). However, when there exists a reserve price, if the maximum bid lies in the interval (v_0, v^R), the seller would keep the object, even when there is a bidder with a higher valuation.

2.3.3 Entry Fees

A reserve price $v^R > v_0$ excludes all bidders with valuations lower than v^R from the auction. Those with valuations exactly equal to v^R are indifferent to participate, because would obtain a zero payoff. An alternative tool that a seller can implement to exclude buyers with low valuations is the collection of an entry fee (i.e. a fixed and non-refundable amount of money that each bidder has to pay, in order to be allowed to participate in the auction).
The same set of bidders that are excluded by a reserve price of v^R can be excluded by setting an entry fee of

$$ef = v^R H(v^R) + \int_0^{v^R} xh(x)dx = \int_0^{v^R} H(x)dx$$

which corresponds to the expected payoff of a bidder with a valuation of v^R, in either format. Clearly, a bidder i with valuation $v_i < v^R$ would not agree to pay the entry fee, ef, in order to participate in the auction. Hence, a reserve price and an entry fee are identifiable, in the sense that for any of them the counterpart tool excludes the same set of bidders.
Many other issues in the static-mechanism approach of auctions have been investigated in literature, including asymmetries in the valuations, interdependences, budget constraints, and changes in risk profiles. However, the core of theory is contained in the main results presented here, and actually it constitutes the baseline model that the new generation of works try to adaptate to dynamic mechanisms.

Chapter 3

Dynamic Mechanism Design

Dynamic or online mechanisms extend the methods of static mechanism design to environments with multiple agents and private information. Decisions must be made as information about types is revealed online and without knowledge of the future. The relevance of such kind of mechanisms is evident since they are pervasive in many environments, such as the following:

i) Selling airline tickets to buyers arriving over time.

ii) Allocating computational resources (bandwidth or CPU) to jobs arriving over time.

iii) Allocating the electromagnetic space of a country to bids arriving online.

iv) Selling a player's contract to another team according to a stochastic stream of bids.

v) Selling adverts on a search engine to a possibly changing group of buyers and with uncertainty about the future supply of search terms.

vi) Allocating tasks to a dynamically changing team of agents.

In each of these settings at least one of the following is true: either agents are dynamically arriving or departing, or there is uncertainty about the set of feasible decisions in the future. The dynamics of agent arrivals and departures, coupled perhaps with uncertainty about the state of the environment, makes the problem of online mechanism design very different from its static counterpart. New considerations include:

i) Decision must be made without information about agent types not yet arrived, along with uncertainty about which decisions will be feasible in future periods.

ii) Agents can misrepresent their arrival and departure times in addition to their valuation for sequences of decisions. Because of this, agent strategies also have a temporal aspect.

iii) Only limited misreports of type may be available. For example, it may be impossible for an agent to report an earlier arrival time than its true arrival.

Last examples show that situtations in which individuals behave strategically under dynamic information disclosure are now widespread, and many aspects of statis mechanism design should adjusted to address such kind of things. A example that clarifies the use of the different pieces of the theory in a dynamic environment is the so-called *secretary problem*. The next section introduces the problem and explores many possible extensions by treating it as an online auction.

3.1 Online Auctions as an Extension of the Secretary Problem

The classical secretary problem can be summarized as follows. A known number n of secretaries are intended to be hired by a manager of a firm, who interview them in an random order within a time interval $\mathcal{T} := [0, T]$, for T fixed and sufficiently large. All secretaries arrive online-uniformly in $\mathcal{T}$, so that all permutations $n!$ are equally likely to be considered. After the manager interviews a secretary i, she is able to learn its quality v_i (presumably a measure of her abilities to perform the job vacant) and rank it with respect to all qualities v_j, $j = 1 \cdots (i-1)$ of candidates interviewed so far. At that point, the manager should make the irrevocable decision of hiring her (in which case the process stops), or rejecting her -interviewing the next secretary in the sequence- and facing the same choice problem as before. It is assumed that the manager incurs in an infinite monetary loss if the vacant is not filled, so if she happens to interview the last secretary in the sequence, then she must hire her.

The manager's aim is to maximize the probability of hiring the best secretary from the pool of applicants, for what she has to assess the tradeoff between stop too soon and wait too much. Since the manager is unable to go back and hire a candidate previously dismissed, if she stops too soon, then incurs in the risk of hiring an apparently-suitable candidate when in fact an even better one might be still to come. On the other hand, if she waits too much, then incurs in the jeopardy of finding out that the best (or even the ones in the top of the distribution) was (were) rejected earlier on.

In the standard formulation of the problem, the candidates do not control the order in which they are interviewed, rather such order is uniformly random[1]. Therefore, if the manager selects a

[1] In a general setting, the order and values v_i could be decided by an adversary, as long as from the perspective

candidate at random, the probability of choosing the best is $1/n$. That unconditional probability actually sets a lower bound for the performance of any other algorithm that takes into account the sequential nature of the problem. In that sense, consider the following strategy suggested to the manager. Interview the first $n/2$ secretaries but do not hire anyone; set $\max_{n/2}(v)$ as the maximum quality encountered in this first half of the pool; and hire the first candidate in the second-half whose quality surpass $\max_{n/2}(v)$. Pursuing such algorithm gives a 1/3-competitive ratio, which is better than choosing at random for all $n > 3$. Lindley (1961) and Dynkin (1963) showed that the optimal generalization of such strategy calls for letting pass the first $k = \lfloor n/e \rfloor$ applicants, and hiring the first one in the complementary half-space whose quality surpass $\max_{n/e}(v)$. This algorithm gives a $1/e$ ratio, which is the most competitive. Actually, such strategy is known as the 37% rule, because the probability of selecting the best candidate is no worse than $1/e \approx 0.37$.

Since it was first introduced, the secretary problem has become a natural starting point in the formulation of more complex optimal stopping problems, and its applications are recurrent in engineering, computer sciences, finance and economics. For example, in the latter field it has been utilized to set optimal contracts in dynamic environments, or as a dual formulation to search theory for finding an optimal job, a suitable partner to trade with, or even, the best spouse to get married. This myriad of applications has conferred it many alternative names in literature, such as the marriage problem, the sultan's dowry problem or the Gogol game.

3.1.1 K-choice and Knapsack Secretary Problems

There have been many endeavors in literature to extend the secretary problem to a broader family of optimal stopping problems. One branch has generated theory and algorithms to solve combinatorial optimization problems such as the k-choice and the knapsack secretary problems. In those, the idea is finding an optimal way to sequentially choose k elements from a set of n, over a finite time interval, to maximize some value function $g(\cdot)$, given a capacity constraint.

The k-secretary problem is a direct extension of the standard problem, ($k = 1$), stated above. Following the same rule as before, Babbaiof et al. (2007) has shown, *via* complex counting, that this algorithm has ratio no worse than e. Nonetheless, Kleinberg (2005) has demonstrated that when k goes to infinity, the corresponding ratio is bounded from below and above by $c + \sqrt{k}$ and $C + \sqrt{k}$, for some constants $0 < c < C$.

On its part, the knapsack problem is a celebrated puzzle in combinatorial optimization theory.

of the manager the probability of facing any sequence still being symmetric.

Its main idea can be captured in the following example. Suppose the government of your country is trying to allocate a given portion $\mathcal{S}$ of the national electromagnetic space among n different competitors, who decide what fraction they want to buy and how much they are willing to pay for it. Each potential buyer i arrive at the same time and place her offer (which is composed by a space request and a maximum willingness to pay for it (s_i, v_i), $0 < s_i \leq S$, $v_i > 0$). If the office in charge of granting the licenses wants to maximize benefits, it has to select which offers accept to maximize $\sum_i v_i$, subject to the capacity constraint $\sum_i s_i \leq \mathcal{S}$. Formally, the ordered pair (s_i, v_i) constitutes a long term contract in which the government commits to allow the firm i to exploit the portion of the electromagnetic space granted for the life of the contract, in exchange for v_i. Notice that a complete solution of the problem requires the verification of all feasible combinations in the value function, which, when n is sufficiently large, cannot be solved in polynomial time (i.e. the problem is non-polynomial complete). However, it is possible to obtain quasi-optimal solutions in reasonable time using meta-heuristics techniques such as simulated annealing, tabu search or genetic algorithms.

When the n potential buyers do not arrive at the same time, but online according to some stochastic process, the problem is called the knapsack secretary problem. In this, the manager does not know the joint distribution of the requests, and therefore when face an offer (s_i, v_i) should evaluate not only the risk that another individual j could offer later a higher value $v_j > v_i$ for the same fraction of the space s_i, but also the risk that very attractive offers could not be accommodated later on if there is not enough room for them. Although the knapsack secretary problem is a harder refinement of the baseline case, the strategy followed to approximate a solution shares the same basic principle of letting pass some offers without accepting any of them, and construct a decision rule based on the information learned. Specifically, the algorithm observes a constant fraction of the offers without accepting any, ranks them by their value-density v_i/s_i, and sets a threshold derived from those sampled. After the initial sequence of offers is considered, elements that surpass the threshold previously defined are chosen. Nevertheless, some "modifications" should be added in order to assure enough room for offers with extraordinary high value-density that arrive in the non-trade segment of the time interval (Babbaiof et al. 2007, 2008).

3.1.2 Time discounting

So far, the kind of problems considered here has been framed in a *time-independent* environment, in which the value attached to an element is independent of the time it was selected. However, this assumption can be overly simplistic in inter-temporal choice problems, especially of the kind widespread in economic applications. Following the electromagnetic space example above, suppose that the office in charge is interested in granting the licenses for perpetuity since the contract is signed, in exchange for some annual fee, v_i. In the notation introduced before, let (s_i, v_i) stand for the space quota and the annual fee the buyer is proposing to the office, and denote $\tau_i = \tau((s_i, v_i))$ the arrival time of (s_i, v_i) in the random sequence of offers. Therefore, if the office accept the request of buyer i, the benefit obtained would be $\pi_i = \sum_{t=\tau_i}^{\infty} v_i d(t)$, where $d(t)$ is an arbitrary decreasing discount time function, known beforehand by all participants[2]. The latter is a version of the knapsack secretary problem with arbitrary discounting, but same principle can be applied to the simple case. In the same fashion, suppose that all secretaries have a constant productivity and infinite life-span. Assume that the firm can obtain a level of benefits v_i if secretary i is hired, so that the present value of the flow of benefits is given by $\pi_i = \sum_{t=\tau(i)}^{\infty}(1/1+r)^t v_i$, where $d(t) = (1/1+r)^t$ is a geometric discount rate for a fixed $r \in \mathbb{R}_+$, and τ_i is defined as before. The use of $1/(1+r)$, where r is the current interest rate in the market, is the most utilized discount factor in economics because it is dynamically consistent[3] in the solution of dynamic programming optimization problems. Moreover, its use improve dramatically the accuracy of the kind of algorithms presented here, which in the case of arbitrary discount rate $d(t)$ is not more competitive that $\sqrt{2}$, even if the entire space of values is known in advance (Babbaiof et al., 2008).

3.1.3 Incentive Compatibility and Incentive Rationality

Up to now, the motivating examples presented to generalize the secretary problem have taken for granted that all agents will be willing to participate in the online mechanism proposed, and will reveal their true valuation to the correspondent decision-maker. However, there might be scenarios in which the design of the online scheme would make beneficial to an arbitrary agent i, who arrived at time τ_i, reveal a different value to v_i. The former are known as incentive-compatible or truthful online mechanisms, which are the case when the objective function $g(\cdot)$ of the decision-maker is monotonic in the one-time value reported by agents. To see this, recall that the optimal strategy for online mechanisms presented so far, asks for letting pass a constant fraction of the offers to extract some information, construct a threshold based on that information, and select the first next element in the sequence that surpasses such threshold.

[2]Although the problem is stated in discrete time for simplicity, its extension to continuous time is immediate.
[3]Hyperbolic discount rates does not share this property for example

According to this rule, there is no incentive for an agent to not reveal her true value, because if she arrives online in the segment at which the decision maker is not accepting offers, is irrelevant what offer she places, but if she arrives in the complementary half-space, then by revealing her true value she maximizes her probability of being selected[4].

Nonetheless, there are situations in which the true revelation discussion is not even a problem because agents would find no rationality to participate in the proposed scheme. When it is the case, mechanisms are called non incentive rational. Online single-object auction with entrant fees constitutes a good example of this sort. In those, for all buyers with an object-valuation below the entrant fee, is irrational to participate in the auction, independently of the auction's format run by the seller. Notice that incentive rationality is not a necessary nor sufficient condition to have a truthful mechanism, and actually there are many examples of truthful mechanisms which are not incentive rational and *vice-versa*.

3.1.4 Endogenous Arrival Time and Revisiting

The reasoning conducted here assume, although, that agents arrive exogenously and uniformly in time, and consequently, that cannot control the moment at which to place their one-time offers. Nevertheless, there are plenty of applications in which agents can endogenously decide the time at which arrive to the process. As a matter of example, consider buyers trying to get airline tickets online, or a commercial bank attempting to buy Britain pounds in the electronic market of a given day. In the first case, a buyer shopping online for airline tickets may choose not to purchase them as soon as the company announces tickets are available, and rather wait until a later date, if they hope prices will go down. In the same fashion, the commercial bank would prefer to buy as soon as possible, if it expects an increase in the exchange rate of the day. When agents can arrive endogenously, sellers shall form a belief about a reasonable distribution of the arrivals in the $\mathcal{T}$.

The endogenous arriving scheme can be even more complicated when there exists the possibility of revisiting. That is, when agents can place more than one offer for the time the online mechanism lasts. Consider a central bank selling a bond, *via* an online mechanism, to a fixed number of bidders n. Assume that each bidder can place more than one offer, but that the bank can identify each offer with the corresponding bidder. Under this scenario, the seller should adjust the threshold defined in the optimal rule discussed in this section, to capture the

[4]Notice that if the secretary is not hired, she will not receive any salary, and therefore she strictly prefers to be hired.

strategic behavior of potential buyers. Observe that a buyer might have an incentive to under declare her true valuation the first time she visits the seller. If the seller is in the stage of getting a feeling of the market, then by under declaring she is contributing to bringing down the cutoff at which the seller might deem convenient to accept an offer. On the other hand, if her offer arrives in the "susceptible to trade" half time space, then buyer might want to shade the bank, in order to get the bond at a lower price than her valuation. However, a buyer conducting such strategy incurs in the risk that her immediate successor reports a value between her true valuation and her last actual value reported, in which case, the process stops and the buyer receives a zero-payoff. A mechanism of this sort, is clearly non-truthful, but still being incentive-rational, due to it is optimal for all agents participate in the process, when an opportunity comes to them.

3.1.5 Unknown Number of Participants

Finally, there have been many developments in literature to extend the secretary problem to applications in which the number of participants is unknown. An illustration, in which this assumption is typical, comes from the plethora of mechanism designed for searching and matching people online. This problem is also called the spouse problem in search theory, and shares the same principles with the pure secretary problem, but with the difference that the seeker does not know the measure of the potential universe which she is dealing with. Notice that as before, the same kind of too early and too late stooping problems are present here, but with the remarkable difference that the decision maker does not have a notion of what too early and too much stand for. The solution proposed in literature takes the number of participants as a random variable over which the decision maker has to hypothesize. Abdel-Hamid et al. (1982) showed there exists an exhaustive family of Bayesian acceptable solutions for a single improper prior distribution.

The so-called *secretary problem* and its different versions, are just one example of a general class of optimal-stopping problems in dynamic environments. There, generally a planner or seller has to decide when stop a selling mechanism in order to maximize revenue. However, such decision rule has to take into account that buyers act strategically, possess private information, and that many times outside options are available. These, and many other theoretical aspects of online or dynamic mechanisms, have been studied in recent literature, stressing properties of efficient implementation, budget constraints, learning over time, and information asymmetries. The next chapter pretends to review (in not an exhaustive way) the recent literature that deals with dynamic mechanisms governed by an embedded optimal-stopping rule.

Chapter 4

Dynamic Mechanisms: Review of Literature

The literature on dynamic mechanism design can be broadly divided into two branches: on one side we can find authors interested in the intricancies of efficient implementation in one-dimmension dynamic models, while on the other hand we can find works dealing with extensions of the uncertainty space in the temporal dimmension, and how it impacts buyers and sellers' equilibrium strategies. The first class of works are mainly focused on prove the existence of Vickrey-Grooves-Clark mechanisms in dynamic environments, and characterize the properties of the payment and allocation rules. In that sense, Bergemann and Valimaki (2010) constructs a payment scheme where each player receives as his payment, her marginal contribution to the social welfare in every conceivable continuation game. In their model, the planner calculates the efficient allocation given agents' reported types, and then also calculates, for each agent i, the optimal allocation when such agent i is excluded from the mechanism. The total expected discounted payment of each agent is set equal to the externality cost imposed on the other agents in the model. In that environment, the social objective is simply to maximize the expected discounted sum of the individual utilities, whose solution is by construction time-consistent.

On its part, Athey and Segal (2007) constructs an efficient, incentive compatible, dynamic mechanism with balanced budgets under a "private values" assumption, wehere agent's payoff does not depend on the other agents' private information, and an "independent values" assumption, in which agent's private signals are independent of other agents' private information, conditional on past public decisions. The main result shows that under these two assumptions, the additional incentive constraints do not rule out the implementation of efficient decision plans using

budget balanced transfers. The desirability of budget balances in the design of a mechanism is well known, since in many times a planner needs to embrace a project without the possibility to count with external funds, even in a dynamic environment. The result by Athey and Segal show that it is possible construct mechanisms of this sort in which agents truthfully reveal their types over time, and agree to follow the allocation and payment schemes proposed by the planner.
Another example of efficient implementation is provided by Mierendorff (2010). In a standard framework of independent private values, where buyers are long lived and arrive randomly over time, the author show that the efficient implementation rule in an online auction of a single object, can be implemented by a mechanism with a *simple payment rule* that generalizes the static Vickrey Auction. The payment scheme derived in the model satisfies the following convenient properties: i) only the winner makes a payment, ii) payments are ex-post individually rational, iii) the mechanism never transfers money to any buyer, and iv) payments are made online (i.e. all information that is needed to determine the payment must be available at the time of allocation). The central result of the paper is that, despite the intertemporal elments in buyers' private information, only one-dimmension is relevant to determine the efficient allocation. For each type, there is a unique period in which a buyer can possibly win the object, and hence only the valuation for this period matters for the efficient allocation rule.
Finally, Pavan et al. (2008) derives a so-called *dynamic payoff formula* which represents the derivative of an agent's expected payoff in an incentive-compatible mechanism with respect to his private information. It provides an envelope-theorem condition summarizing local incentive compatibility constraints. Intuitively, such formula represents the impact of an infinitesimal change in the agent's current type on agents' equilibrium expected payoff. According to the authors, this change can be decomposed into two parts, being the first the familiar effect of the current type on agent's expected utility, and the second the indirect effect of the current type on the expected utility through its impact on the type distributions in each of the subsequent periods. Necessity and sufficiency conditions on the *dynamic payoff formula*, of a fairly general class of incentive compatible mechanisms, are also derived.
Other side of the literature has taken the road of relaxing the conventional one-dimmension uncertainty in agents' space, to include explicitly arrival and departure times as part of their private information set. This extension brings new considerations on the strategic behavior of buyers, as well as on the conditions that assure incentive compatibility and incentive rationality. Moreover, it permits to tackle new interesting issues such as incorporating learning capabilities by the seller and or buyers.

In this stream, Pai and Vohra (2008) consider a dynamic auction problem in which a seller endowed with C units of an object faces potential buyers with unit demand who arrive and depart over T periods. Buyers' type-space is three-dimmensional, composed by their arrival and departure time, as well as by their private valuation over the object. In such setting, authors derive the revenue maximizing Bayesian incentive compatible selling mechanism. The approach followed is an extension of Myerson (1981) paper on optimal auction design in static mechanisms, however the results found do not follow immediately, since now incentive compatibility shall to be verified for all the directions in the type-space. In a similar fashion, Mierendorff (2010) explores a model to sale an object in a dynamic environment where agents have the same type-space as before, but where arrival times are public information for the seller, and agents' degree of patience is given by exogenous deadline times. When agents participate in the mechanism, they are prompted by the seller to declare their private valuation as well as their deadlines. The author first explore the optimal mechanism neglecting the incentive constraint for the deadline, and then analyze conditions that could make the relaxed solution suboptimal. That is, situations that lead violation to the neglected incentive constraint. Unlike Mierendorff, Galiien (2005) considers a model with only patient buyers that have to report arrival times and private valuations. In an infinite horizon model it is shown that under certain conditions on the inter arrival time distribution, the optimal mechanism is the same as with impatient or non strategic buyers.

In general, these models assume independence between the arrival time and the valuation of the buyers, nevertheless it might be the case they are correlated, which changes the conditions on the payment scheme that assures incentive compatibility. Indeed, the efficient allocation rule cannot always be implemented following the four criteria pointed by Mierendorff (2009), if independence is not satisfied. Moreover, since buyer's can only delay but not bringing forward their arrivals, the mechanism could pay a subsidy in order to induce agents report earlier their arrivals, with the purpose to take advantage of the externality in information. Notice, though that the subsidy shall be paid to all buyer's and not only to the winner, which assumes the seller has external resources to do so.

The present paper, like Mierendorff (2010), assumes arrival times are observable by the seller, but instead of considering exogenous deadlines, allow buyers to optimally select their departure time under the presence of external trading opportunities. Incentive compatibility and implementation issues are explored in such environment, where additionally time is partitioned in disjoint market windows.

Chapter 5

An optimal stopping dynamic mechanism under market windows and endogenous departure times

A risk-neutral, revenue-maximizing seller endowed with an indivisible, storable object, faces a continuous stream of self-interested, risk-neutral potential buyers, who arrive in an interval $\mathcal{T} \equiv [0, T]$ (which corresponds to the useful-life of the object). As soon as the seller is endowed with the object, she advertises a sale opportunity that opens the market immediately. Nonetheless, trading opportunities are constrained by the legal and institutional framework of the economy, that defines official periods at which exchange can take place. The set of these periods is denoted by W and each of its elements are known as market windows. Consequently, $\mathcal{W} = W \cap \mathcal{T}$ is the set of authorized periods at which a deal over the object can occur in the model. Market windows are assumed to be truncated and symmetric in $\mathcal{T}$, and without loss of generality, the earliest window starts at $t = 0$[1].

The set of agents is denoted by $\mathbb{I} = \{1, 2, \cdots I, I+1\}$, with the convention that $i \leq I$ represents a bidder and I+1 represents the seller. Agents' arrivals in each window are described by the stochastic process $\{\mathcal{N}(t), t \in w_k\}$, where $\mathcal{N}(t)$ is a random variable representing the number of arrivals up to time $t \in w_k$. If $t \in (\overline{t_k}, \underline{t_{k+1}})$, $\mathcal{N}(t) = 0$ for all k. That is, *inter-window periods* are considered dead times in which no potential buyer arrives to place a bid. Notice however, that in a model with commitment on the side of the seller, it might be rational for some potential buyers

[1]According to this description, for all uniform partition $\{t_0 = 0 < t_1 < \cdots < t_n = T\}$, $\mathcal{W} := \{w_j : w_{j+1} = [t_{2j}, t_{2j+1}], \ j = 0, 1, \cdots (n-1)/2\}$. Relabeling bounds, we will refer of a window w_k as the interval $[\underline{t_k}, \overline{t_k}]$

to arrive in dead times to arrange an option for buying in the next market window. Nonetheless, such cases are rule out in the present model, and instead, bids, payments and allocations only occur within market windows.

Seller's uncertainty about the value estimate of bidder i can be described by a continuous probability distribution over a finite interval. In that sense, before getting into a market window, each of the I potential buyers privately and independently learns her actual valuation[2] from the compact support $V \equiv [\underline{v}, \overline{v}] \in \mathbb{R}_+$, where $-\infty < \underline{v} < \overline{v} < +\infty$[3], according with a continuous probability density function $f_{w_k}(\cdot)$, which satisfies $f_{w_k}(v) > 0$ for all $v \in V$. As usual, $F_{w_k}(\cdot)$, will denote the cumulative distribution function corresponding to $f_{w_k}(\cdot)$.

Actual valuation of each buyer i might depend on the history of past valuations. Specifically, it could be assumed that $v_i(w_k)$ is drawn from a history-dependent probability distribution. $F_{w_k}(\cdot|v_i^{w_{k-1}})$ on V, where $v_i^{w_{k-1}}$ represents buyer i's history of valuations up to window w_{k-1}. Generally, F_{w_k} is an element of $\mathfrak{F} \equiv \{F : \mathcal{B} \rightarrow [0,1]\}$, the set of all possible measures on V. Nonetheless, in order to maintain tractability independence of distribution along market windows is also assumed.

Moreover, following Myerson (1981) the next regularity condition is imposed.

Assumption 1. *For all $w_k \in \mathcal{W}$, $C(v) = v - \frac{1-F_{w_k}(v)}{f_{w_k}(v)}$ is strictly increasing in v.*

$C(\cdot)$ is called the virtual valuation of a buyer given her actual valuation v, and its well defined in virtue of the full support assumption of f_{w_k}. As we will see later, this condition assures that for every pair of bidders i and j, such that $v_i > v_j$, bidder i will have a higher probability to get the object.

Along with her valuation $v_i(w_k)$, and arrival time $a_i(w_k)$, a type of an agent is completed by her claimed departure time of the current market window, $d_i(w_k)$, which stands for the time up to which the buyer is willing to keep her offer active in the current window. For convenience, independence between arrival times and valuations is assumed for all buyers along all windows. Hence, agent's multidimensional full type within market window w_k corresponds to:

$$\Theta_i(w_k) := \{\theta_i(w_k) \equiv (a_i(w_k), d_i(w_k), v_i(w_k)) \in [\underline{t_k}, \overline{t_k}) \times (\underline{t_k}, \overline{t_k}] \times V : a_i(w_k) < d_i(w_k)\} \quad (5.1)$$

Its natural restriction to time t, given by $\Theta_i^t(wk)$, additionally requires that $a_i(w_k) < t$.

[2]Those valuations might change along market windows to reflect the fact that preferences toward the object can be affected by random events during its useful life, or might be revised on the light of new information conveyed by the behavior of competitors in previous windows. In other words, and contrary to other dynamic mechanisms as Said (2010), buyers' valuations are not persistent.

[3]The space of valuations V, along with the *Borel* sigma-algebra $\mathcal{B} \in \mathbb{R}$ conform a measurable space.

Many papers such as Gallien (2006), Pai and Vohra (2008), and Gershkov and Moldovanu (2010) consider buyers' arrival times as private information. However, for the sake to isolate the intricancies related with the report of departure times, this paper assumes arrivals are always observable to the seller. A similar characterization of agents' types is given in Meriendorff (2010). However, in his model deadlines are given exogenously (as a reflection of the heterogeneity in buyer's degree of patience), whereas in the present model agents have to choose an optimal departure time report at their arrival. In fact, all potential buyers are assumed to be patient/strategic, whose actual departure time of the mechanism depends on the availability of trading opportunities in parallel markets, as in Rochet and Stole (2002).

More precisely, once an agent i enter into a market window, she receives -independently and exclusively- a private signal about a bilateral trading opportunity[4] over a substitute object for which buyer i assigns a valuation $\tilde{v}_i(w_k)$: a realization from the conditional distribution $F_{w_k}(\cdot|\tilde{v}_i < v_i)$. The time at which the signal arrives in any window is distributed according to $G_{w_k}(\cdot) \equiv G(\cdot)$, and correspondent continuous density $g(\cdot)$, with the property that $g(t) > 0$ for all $t \in w_k$[5]. Moreover, in the bilateral bargaining process the buyer almost extracts the whole surplus, in the sense that the negotiated price is fixed at $\epsilon > 0$, for ϵ sufficiently small. When the signal informs about the alternative market, buyers have to make an instantaneous decision of leaving the actual market window, or stay there up to her initially claimed departure time. Nonetheless, once the possibility of leaving the actual market is rejected, it cannot be recalled. To complete the environment, we assume that if buyer i reports a departure time $\hat{d}_i(w_k)$, but effectively leaves at $d_i(w_k) < \hat{d}_i(w_k)$, and the seller happens to allocate her the object at $\tau_{w_k} \in (d_i(w_k), \hat{d}_i(w_k)]$, then agent i does not get the object but has to pay a *non-commitment cost*, which is equal to the elapsed time between allocation and her actual departure time, multiplied by a penalty fee of $\lambda > 0$. Nonetheless, this penalty fee is only applicable to the winner, though it might be the case that some of the non-allocated buyers have departed earlier as well[6]. Therefore, unlike conventional mechanisms, collection of fines is another potential source of revenue for the seller.

Buyers are assumed to have unit demand in every window, so that if an agent walks away from the institutionalized market at some w_k, it does not preclude her to come back to the market at w_{k+1} to bid for the object if it has not been allocated yet. This feature, along with the fact

[4]The appearance of such opportunities are natural under asymmetries in the information setting, or when agents decide to invest in order to search for better opportunities in less known markets.

[5]As usual in Bayesian games, Θ, F_{w_k}, f_{w_k}, G_{w_k} and g_{w_k} are public information.

[6]This payment scheme is supported by the fact that verify the true state of potential buyers is costly

that agents only discover the bargaining bilateral trading opportunity when they have already arrived to the market, let the stochastic process $\mathcal{N}(\cdot)$, represent the stream of agents at any window. For tractability, buyers are allowed to place just one bid in any market window, which has to be reported along with her departure time at their arrival -those cannot be updated once revealed. It is important to remark that seller cannot see whether a potential buyer have walked away before her reported departure time, and only can verify this condition if decides to allocate her the object before such time.

Unlike other online mechanisms where if an offer is not attended immediately then can never be considered again[7], here the possibility of recalling is considered -in the spirit of Zuckerman (1986). However, it is qualified to the current active window, that is, an allocation at time $\tau_{w_k} \in w_k$ for any w_k, can be and only can be mapped to any agent i with arrival time $a_i(w_k) < \tau_{w_k}$.

With respect to seller's valuation, we suppose at any time she attaches a v_0 value to the object at hand, which is the utility she can derive if decides to consume the good (or execute it in the case of a financial asset), instead of allocating it. Observe that if the seller conserves the object after w_k window's closing, then she necessarily has to keep it until w_{k+1} window's opening if she wants to allocate it to someone else. Once the seller allocates or consumes the object, the market is closed and resale is not allowed. Thus, the seller only can allocate the object within a market window, but can consume it at any time $t \in [0, \mathcal{T}]$. Like potential buyers, seller discounts time using the factor $\delta \in (0, 1)$.

From seller's perspective, we can impose a partial order $\succeq$ on $\Theta(w_k)$ for all w_k -as in Pai and Vohra (2008)- such that

$$\theta_i(w_k) \succeq \theta_j(w_k) \equiv (a_i(w_k) \leq a_j(w_k)) \wedge (d_i(w_k) \geq d_j(w_k)) \wedge (v_i(w_k) \geq v_j(w_k)) \tag{5.2}$$

Intuitively, a type will be weakly preferred to another if the buyer to which it is referred arrives earlier, reports a later departure time, and has a weakly higher valuation. Suppose we have types $\theta_i(w_k)$ and $\theta_j(w_k)$ such that $a_i(w_k) = a_j(w_k)$, $d_i(w_k) < d_j(w_k)$, and $v_i(w_k) > v_j(w_k)$. Then, we can interpret the ex ante seller's "marginal rate of substitution" between valuations and departure times, as the tradeoff of having the possibility to allocate the object to the buyer with higher valuation $v_i(w_k)$, up to her claimed departure time $d_i(w_k)$, versus the option of extending the possible optimal allocation up to $d_j(w_k)$. In this case, seller can expect the arrival of a buyer with type $\theta_l(w_k)$, such that $a_l(w_k) \in (d_i(w_k), d_j(w_k))$, and $v_l(w_k) > v_j(w_k)$. Clearly, expected value $v_l(w_k)$ should be high enough to compensate for the loss of value in time.

[7] Think for example in the secretary problem and its many extensions, as presented in Babaioff et al. (2006).

Agents' utility is quasi-linear, which implies that if a buyer i arrives at time $a_i(w_k)$, gets the object at $\tau_{w_k} \in [a_i(w_k), \overline{t_k}]$, and pays $p_i(w_k)$, her utility is given by $e^{-\delta\tau_{w_k}}(v_i(w_k) - p_i(w_k))$. Correspondingly, the utility of the seller would be $e^{\delta\tau_{w_k}}(p_i(w_k) - v_0)$. Notice the attention is restricted to payment schemes where full transfer from buyers occur instantaneously upon the allocation of the object[8]. Furthermore, no budget constraint is imposed to potential buyers, meaning that at any time they have enough resources to pay for their bids, and so never incurred in default with the seller.

5.1 Allocation structure

Since recall may be utilized by the optimal policy, the state of the game can be summarized by the history of types reported by agents throughout time. Borrowing the general setting of Gershkov et al. (2010), let

$$\mathcal{H}_{w_k}^t = \prod_{\mathcal{N}(t)=0}^{I} \Theta^t(w_k)^{\mathcal{N}(t)} \tag{5.3}$$

be the set of all possible histories, $h_{w_k}^t$, at time t within the active market window w_k. Here, we set $\Theta_i^t(w_k) = \Theta^t(w_k)$ since buyers are symmetric.

For convenience in notation, let $h_{w_k}^t = (h_i^t(w_k), h_{-i}^t(w_k))$, where $h_i^t(w_k)$ denotes the reported type of agent i up to time t in window w_k (with the convention it is the empty set if agent i has not arrived at time t[9].); and $h_{-i}^t(w_k)$ denotes the derived history of types up to time t, reported by potential buyers other than i. In addition, define $\mathcal{S}_{w_k}^t := \{s_{w_k}^t : s_{w_k}^t = (h_{w_k}^t, z_{w_k}^t) \in \mathcal{H}_{w_k}^t \times \mathbb{I} \cup \{0\}\}$, where $z_{w_k}^t$ indicates if the object is still available at time t within window w_k, $z_{w_k}^t = 0$, or has already been allocated to some agent i, $z_{w_k}^t = i$.

From the latter, we can build $\mathcal{H}_{w_k}$, the space of histories in any active window w_k, as the continuous product of $\mathcal{H}_{w_k}^t$ in $[\underline{t_k}, \overline{t_k}]$. To conform its correspondent measure space we appeal to Kolmogorov's extension theorem[10] to let $\mathcal{F}_{w_k}$ be the σ-algebra generated by $\mathcal{H}_{w_k}$, μ_{w_k} its induced measure, and $\{\mathcal{F}_{w_k,t}\}_{t \geq \underline{t_k}}$ its associated filtration. As above we set $h_{w_k} \in \mathcal{H}_{w_k}$ as $(h_i(w_k), h_{-i}(w_k))$. With these fundamentals, the time at which seller stops the market is defined

[8] Nevertheless, differences in allocation and payment times are customary in markets where a deposit in advance is necessary to meet the institutional requirements that allow the buyer to take possession of the object. In fact, any market in which a guarantee deposit is necessary, like electromagnetic space auctions or public project tenders, is classified in this category.

[9] If agent i arrives to the mechanism, then given that buyers can place only one bid, clearly $h_{w_k}^t = \theta_i^t(w_k)$.

[10] For a detailed explanation of the theorem as well as the description of the two consistency conditions embedded see Oksendal (2002).

as following:

Definition 15. *A stopping time with respect to* $\{\mathcal{F}_{w_k,t}\}_{t\geq \underline{t_k}}$*, in the current window* w_k*, is defined as a function* $\tau_{w_k} : \mathcal{H}_{w_k} \to [\underline{t_k}, +\infty)$*, where* $\{h_{w_k} \in \mathcal{H}_{w_k}; \tau_{w_k}(h_{w_k}) \leq t \in \mathcal{F}_{w_k,t},\ \text{for all}\, t \geq \underline{t_k}\}$.

Clearly, the optimal stopping time within window w_k can be also defined as $\tau_{w_k} \equiv \inf\{t \in [\underline{t_k}, \overline{t_k}] : z^t_{w_k} > 0\}$, in virtue of our notation. Denote $\tau_{w_k}(\theta_i(w_k), h_{-i}(w_k))$ the optimal stopping time within window w_k if agent i reports a type of $\theta_i(w_k)$, while the rest $I-1$ buyers conform a history $h_{-i}(w_k)$. If the object is not allocated within the current window, then we set $\tau(h_{w_k}) = \infty$. Consequently, in all subsequent windows to allocation or unilateral closing, stopping time is equalized to infinity.

The subset of histories $h_{-i}(w_k)$ in which the agent i gets the object, when reporting a type $\theta_i(w_k)$, is denoted by $\mathcal{H}_{-i}(\theta_i(w_k))$, while $\mathcal{H}^c_{-i}(\theta_i(w_k))$ denotes its complement. Furthermore, to keep track of the relevant histories for the mechanism, let $A(s^t_{w_k}) \in \mathbb{I}$ be the set of active buyers in state $s^t_{w_k}$, or in other words, the set of agents that have arrived and have not departed from the mechanism at time t in the current window w_k (i.e. those whose initially-reported departure time has been not surpassed, and have not walked away at time t). Naturally, no agent is active if the market is closed.

Observe that unlike standard stopping-time problems, even if histories coincide in arrivals and reported valuations, the optimal stopping time might depend on the set of active buyers in any window.

5.2 Direct mechanisms

Under this environment, seller aims to design a mechanism whose equilibrium maximizes seller's expected revenue. Nonetheless, the degree of generality embedded in the environment stated above is a major difficulty to find a solution to its correspondent optimality problem, because it is necessary to consider any class of mechanism. However, as it is typical in such studies, this difficulty is resolved by appealing to the *Revelation principle*, a result firstly introduced by Myerson (1981). According to this outcome, it is just necessary to qualify the search to the class of mechanisms in which agents are directly asked to report their types, since any implementable rule associated with an arbitrary strategy space can be supported by the equilibrium solution induced by its correspondent direct mechanism.

An *allocation rule* is a function $y^t_{w_k} : \mathcal{S}^t_{w_k} \to [0,1]^{\mathcal{N}(t)}$ that defines the winning probability of active buyers in state $s^t_{w_k}$ -given $z^t_{w_k} = 0$. Since there is only one object available in the market and resale is not allowed, an allocation rule must satisfy the following feasibility constraint

$$\forall t,\ w_k,\ s^t_{w_k}: \sum_{i\in A(s^t_{w_k})} y_i(s^t_{w_k}) \leq 1 \tag{FC}$$

Likewise, a general *payment rule* is a function $p^t_{w_k} : \mathcal{S}^t_{w_k} \to \mathbb{R}^{\mathcal{N}(t)}$ that sets a price to be collected at any time the market is open. Negative prices stand for transfers from the seller to buyers.
We call an allocation rule discrete, if it allocates the object exclusively upon arrival. In other words, if $\tau_{w_k} < \infty$, then $\tau_{w_k} = a_i(w_k)$ for some i. However, since recall is allowed and buyers are patient, the last expression does not mean that the object shall be allocated to the buyer with contemporary arrival time to the optimal stopping rule; but instead to an active agent j, with $a_j(w_k) < a_i(w_k)$.
Therefore, in this environment a direct mechanism takes any possible history in the game for all market windows, and decides whom to assign the object (if it has not assigned before), and what price charges for it. The following definition formalizes this notion.

Definition 16. *A direct revelation mechanism is a sequence of allocation and a payment policies* $\Gamma \equiv \{(y_{w_k}, p_{w_k})\}_{w_k\in\mathcal{W}}$ *such that: (i)* $(y^t_{w_k}, p^t_{w_k})$ *is* $\mathcal{F}_{w_k,t}$*-measurable, for all* w_k, $t \in w_k$*; (ii)* $y^t_{w_k}$ *satisfy the feasibility constraint (FC); and (iii)* $y_i(s^t_{w_k}) = 0$, $p_i(s^t_{w_k}) = 0$ *for all* $i \notin A(s^t_{w_k})$.

Moreover, the expected winning probability of a buyer who arrives to the mechanism at time $a_i(w_k)$, and reports a type $\theta'_i(w_k) = (a_i(w_k), d'_i(w_k), v'_i(w_k))$ -given that the rest of the potential buyers report truthfully their types- corresponds to

$$\pi(\theta'_i(w_k)) = \mu_{w_k}|a_i(w_k)\{\mathcal{H}_{-i}(\theta'_i(w_k))\} \tag{5.2.5}$$

where $\mu_{w_k}|a_i(w_k)$ is the conditional measure on $\mathcal{H}_{w_k}$ upon arrival of agent i.

5.3 Buyers' optimal departure time report

Recall that in the present environment as soon as the potential buyers get into the market, they learn their valuations and instantaneously have to declare a type[11]. Once buyers are on the market, they are eligible to receive a signal about a parallel bilateral trading opportunity, according to some distribution $G(\cdot)$. In such bilateral bargaining, a potential buyer can obtain

[11]Such type includes a report a of a departure time that cannot be modified once it is claimed.

a surplus of $\tilde{v}_i - \epsilon$ (where $\tilde{v}_i$ is a realization from $F_{w_k}(\cdot|v_i < \tilde{v}_i)$ and $\epsilon \to 0$), if she decides to take the option of walking away. However, there is a cost if the mechanism allocates the object to a buyer but it happened she has departed earlier -the cost being proportional to the time elapsed between the actual departure time and the reported one. The analysis of the tradeoff in reporting a departure time is as follows. A higher reported time $d_i(w_k)$ increases the probability to incur in a *non-commitment cost*, since the probability of the event that the signal arrives in $[a_i(w_k), d_i(w_k)]$, and the individual effectively walks away, increases with $d_i(w_k)$. On the other hand, since the optimal policy allows recalling, when the agent reports she is willing to sustain the offer for longer time, then she increases the probability of being allocated within the institutionalized market. Hence, when declaring a departure time agents seek to maximize the ex-ante value to participate in the market window, but considering the possibility of walking away at any time[12].

$$\tilde{d}_i = \arg\sup_{d_i \in \mathbb{C}_i} \mathbb{E}_{\mathcal{H}_{-i}(\theta_i)}|v_i \left[\int_{a_i}^{\tau_{w_k}} e^{-\delta t}\phi(v_i) - \left(e^{-\delta \tau_{w_k}}[\lambda(d_i - t) + (v_i - p_i)] \right) g(t)dt \right] \pi(\theta_i) \quad (5.3.6)$$
$$+ \int_{a_i}^{t_k} e^{-\delta t}\phi(v_i)g(t)dt[1 - \pi(\theta_i)]$$

Letting $\mathbb{C}_i := [a_i, \overline{t_k}]$, and $\phi(v_i) := \mathbb{E}[\tilde{v}_i - \epsilon|\tilde{v}_i < v_i]$, equation (5.3.6) states the correspondent optimization problem for the agent.

The first term in (5.3.6) accounts for the expected utility of the buyer when she walks away earlier than the allocation time[13]. Her net gain is determined by the discounted expected value of $\tilde{v}_i - \epsilon$, given the conditional density of the signal, minus the potential cost in which the buyer incurs for departing before the time she claimed. Such cost includes not only the so-called *non-commitment* portion, but also the utility lost if the object would have been allocated to her. Observe this cost is discounted using the allocation (rather than the current) time, since it is until such time when the seller discovers the potential buyer has effectively failed in his commitment to buy. Finally, the second term summarizes the utility the agent can get in the bilateral trading, given she is not allocated in the institutionalized market.

The departure time claimed affects the stopping rule in the optimal allocation $\tau_{w_k}(\theta_i, h_{-i})$, as well as the measures over the subsets at which the agent is and is not allocated. In fact, *ceteris*

[12] For convenience, the dependence of variables to the actual window w_k will be omitted henceforth in the absence of ambiguity.

[13] Recall that if the object is allocated within the current window $\tau(h_{w_k}) < \overline{t_k}$, otherwise it is set equal to infinity.

paribus a later departure time increases the measure of $\mathcal{H}_{-i}(\theta_i)$, and consequently decreases the measure of its complement $\mathcal{H}^c_{-i}(\theta_i)$. Notice that even agents are homogeneous, optimal departure times are determined by the correspondent realization of individual valuations, as well as by the arrival times.

5.4 Buyers' instantaneous decision of walking away and expected departure time

Despite the fact that buyers have to *ex ante* declare a departure time, they should decide whether leaving the market or not when have the opportunity to do so. That is, when an on-market buyer with type θ_i^s effectively receives her signal at arbitrary time $t < s < d_i$ over a substitute object with valuation $\tilde{v}$, she has to decide if walking away at that moment, or dismisses that opportunity to conserve the possibility of being allocated in the institutionalized market. So, if $\kappa_i^s := \kappa_i^s(\theta_i, \tilde{v}_i)$ denotes the probability that a buyer i walks away when receiving the signal at time s, given the valuation of the outside opportunity is $\tilde{v}$, then

$$\kappa_i^s(\theta_i, \tilde{v}) = \begin{cases} 1 & \text{if} \qquad e^{-\delta s}(\tilde{v}_i - \epsilon) > \mathbb{E}_{\mathcal{H}_{-i}(\theta_i)}[e^{-\delta d_i}(v_i - p_i + \lambda(d_i - t))]\pi(\theta_i) \\ 0 & \text{otherwise} \end{cases} \tag{5.4.7}$$

In other words, the potential buyer would take the bilateral opportunity with certainty if the realized discounted utility surpasses the right hand side in the latter expression, which in turn is the opportunity cost conformed by the *non commitment* fine and the surplus derived if she would have won the object in the institutionalized market. Therefore, the unconditional probability of walking away at an arbitrary time t is given by, $\rho_i^t = \mathbb{E}_{F_{w_k}(\cdot|\tilde{v}<v)}[\kappa(\theta_i, \tilde{v})]$.
According to the latter, the probability $\alpha_i^t := \alpha_i^t(\theta_i)$ that a buyer actually remains in the mechanism at time t, given she has arrived, is the union of two disjoint events: the event that the signal have not arrived at time t, and the event that the signal have arrived but the agent have decided to dismiss that opportunity to continue in the institutionalized market. Using the notation introduced before, $\alpha_i^t = \int_t^{t_k} g(s)ds + \int_{a_i}^t g(s)(1-\rho_i^s)ds$. From this structure, it is straightforward to build buyers' expected departure time at their arrival as

$$\mathbb{D}_i \equiv \int_{a_i}^{d_i} t\rho_i^t g(t)dt + \alpha_i^{d_i} d_i \tag{5.4.8}$$

The first term in (5.4.8) accounts for the expected time the buyer leaves if attends the signal before the optimal claimed time, while the second one is precisely the probability that she remains in the mechanism up to her initially-reported time.

Proposition 6. *For all $h_{-i} \in \mathcal{H}_{-i}$, define the operator $\chi : \mathbb{C}_i \to \mathbb{C}_i$ such that $(\chi\tilde{d}_i)(a_i, v_i)$ is given by the right hand side of (5.3.6). Then, χ has a fixed point $\hat{d}_i \in \mathbb{C}_i$.*

Proof. Fix $h_{-i} \in \mathcal{H}_{-i}$. Given that seller's preference over types is described by (5.2), for all ordered pairs (a_i, v_i) of an arrival time and a valuation, seller would prefer higher departure times. Now consider the following two cases.

Case 1. If there is no $d_i \in [a_i, \overline{t_k}]$ that prompts the seller to allocate the object to buyer i, then the measure of $\mathcal{H}_{-i}(\theta_i)$ is zero, and consequently her expected net gain would correspond to $\int_{a_i}^{t_k} e^{-\delta t}\phi(v_i)g(t)dt$. Here, the solution of (5.3.6) is the whole interval $[a_i, \overline{t_k}]$. Without loss of generality, buyer reports $d_i = a_i$ and becomes an impatient buyer.

Case 2. Suppose on the contrary that there exists $d_i \in [a_i, \overline{t_k}]$ such that if reported by buyer i -in conjunction with (a_i, v_i)- the seller would allot her the object. In other words, suppose $\tau_{w_k}(\theta_i, h_{-i}) < \overline{t_k}$ and $y_i(s_{w_k}^{\tau_{w_k}}) = 1$. Moreover, if we let $\underline{d_i}$ denote the minimum d_i that satisfies such condition, then

$$\underline{d_i} = \inf\{d_i : d_i > d_k \text{ such that } v_k > v_j, \text{ for all } j \in A(s_{w_k}^{a_i})\} \tag{5.4.9}$$

In (5.4.9), $\underline{d_i}$ represents the minimum time that makes agent i the buyer with the highest valuation in the mechanism, based exclusively on the information collected at his arrival time a_i. Clearly, if $d_i' < \underline{d_i}$ then there exists at least another agent with higher valuation and therefore agent will not be allocated. On the other hand, all d_i's with $d_i' > \underline{d_i}$ have higher *non-commitment* costs, since the signal is distributed according to a full support distribution $G(\cdot)$; and therefore they are inefficient. Notice that if there are new arrivals in period $(a_i, \underline{d_i})$ there are three interesting scenarios. The first one is the arrival of at least one buyer with higher valuation than v_i and later departure time than d_i, in which case buyer i is not allocated. The second one, refers to the situation where all buyers that arrive with higher valuation than v_i declares an earlier departure time. Here, buyer i will have a positive probability of being allotted if seller delays allocation time beyond the later departure time of those buyers with higher valuations than buyer i. However, there might be the case, that buyers with lower valuations arrive and declare later departure times than d_i to have positive probability of being allotted. In this case, agents have to compute the minimum departure time $\underline{\underline{d_i}}$ such that $\int_{a_i}^{\underline{\underline{d_i}}} e^{-\delta t}\phi(v_i)g(t)dt$ equalizes $\int_{a_i}^{\underline{\underline{d_i}}} e^{-\delta \underline{\underline{d_i}}}[\lambda(\underline{\underline{d_i}} - t) + (v_i - p_i)]g(t)dt$. Therefore, in such cases where $\underline{d_i} > a_i$, the optimal reported time for a given history $h_{-i} \in \mathcal{H}$ is given by $\hat{d}_i = \max\{\underline{d_i}, \underline{\underline{d_i}}\}$. ■

On the light of proposition 1, the optimal departure time reported by buyer i at their arrival, can be computed as the expected solutions of (5.3.6), over all possible histories $h_{-i} \in \mathcal{H}_{-i}$, given a valuation v_i. That is,

$$d_i^* = \mathbb{E}_{\mathcal{H}|v_i}[\hat{d}_i] \tag{5.4.10}$$

Corollary 1. *For all buyer i in any window w_k, the optimal departure time report, $d_i^*(w_k)$ is increasing in the realization of buyer's valuation $v_i(w_k)$.*

Proof. See the appendix ■

5.4.1 Incentive compatibility and incentive rationality

Notice that in virtue of proposition 1, the space of types for all agents consists of just one dimension, since their optimal departure time depends on v_i, and agents will not depart from its report. In fact, d_i and v_i are correlated (see proposition 2), and seller can observe if a buyer i is misrepresenting his departure time based on the report of his valuation.
To arise to incentive compatibility constraints of the mechanism, we will first analyze the amount a buyer expects to pay when entering to the mechanism if reports a type $\theta_i^{'} = (a_i, v_i^{'}, d_i^*(v_i^{'}))$[14]. Recall that only active agents are subject of charge and transfers are produce only at allocation time. In that sense the expected payment is given by,

$$\Psi(\theta_i^{'}) = \mathbb{E}_{\mathcal{H}_{-i}(\theta_i^{'})}\left[\int_{a_i}^{d_i^{'}} e^{-\delta d_i^{'}}[\alpha_i^{d_i^{'}} p_i(s_{w_k}^{d_i^{'}}) + (1-\alpha_i^{d_i^{'}})\lambda(d_i^{'} - t)]dt\right]\pi(\theta_i^{'}) \tag{5.4.11}$$
$$+\mathbb{E}_{\mathcal{H}_{-i}^c(\theta_i^{'})}\left[e^{-\delta\tau_{w_k}} p_i(s_{w_k}^{\tau_{w_k}})\right](1-\pi(\theta_i^{'}))$$

The first expectation corresponds to the payment buyer faces when he is allocated, given he reports a type $\theta_i^{'}$ (i.e. the expectation is taken over $\mathcal{H}_{-i}(\theta_i^{'})$ Here, the payment is a convex combination of the payment charged in the institutionalized market and the fine agent has to pay if leaves earlier. Notice that limits of integration reflects the fact that if a buyers is allotted, then it is at his deadline time. The second expectation is taken over subset of histories at which the buyer is not allocated, which is equal to zero in all *only-winners-pay* mechanisms. Henceforth, we will consider only this kind of payment schemes.

[14]For easiness of notation we will abbreviate $\theta_i^{'}$ as $(a_i, v_i^{'}, d_i^{'})$.

Using (5.2.5) and (5.4.11), the expected utility of participating in the mechanism corresponds to,

$$U(\theta_i'|\theta_i) = \mathbb{E}_{\mathcal{H}_{-i}(\theta_i')}\left[e^{-\delta d_i'}v_i - \Psi(\theta_i')\right]\pi(\theta_i') \tag{5.4.12}$$

Expected utility from truth-telling is simply written as $U(\theta_i)$.

Definition 17. *A direct mechanism Γ is Bayesian incentive compatible and incentive rational, if for all $i \in \mathbb{I}$, $w_k \in \mathcal{W}$, and $\theta_i(w_k)$, $\theta_i'(w_k) \in \Theta_i(w_k)$ it holds that:*

$$U(\theta_i(w_k)) \geq U(\theta_i'(w_k)|\theta_i(w_k)) \tag{IC}$$

$$U(\theta_i(w_k)) \geq 0 \tag{IR}$$

Proposition 7. *Let Γ a direct mechanism that allocates only at the deadline. Then Γ is incentive compatible if and only if for all $i \in \mathbb{I}$, $v_i(w_k), v_i'(w_k) \in V$, $a_i(w_k) \in [\underline{t_k}, \overline{t_k}]$, and $d(w_k) \in [a_i(w_k), \overline{t_k}]$:*

$$v_i(w_k) > v'(w_k) \Rightarrow \pi(\theta_i(w_k)) \geq \pi(\theta_i'(w_k)) \tag{IC.1}$$

$$U((a_i(w_k), d_i(v_i(w_k)), v_i(w_k)) = U((a_i(w_k), d(\underline{v}), \underline{v})) + \int_{\underline{v}}^{v_i} e^{-\delta d_i}\pi((a_i(w_k), s, d_i(s)))ds \tag{IC.2}$$

$$U((a_i(w_k), \underline{v}, d(\underline{v}))) \geq 0 \tag{IC.3}$$

Proof. Clearly, (IR) implies (IC.3). Moreover, since the mechanism is incentive rational and incentive compatible, all agents agree to participate and will reveal their true types. Since seller assigns the object to the agent with the highest valuation, a higher valuation implies a higher probability of being allotted. In other words, $\pi(\cdot)$ is monotone.

Now, notice that

$$\begin{aligned} U(\theta_i'|\theta_i) &= \mathbb{E}_{\mathcal{H}_{-i}(\theta_i')})[e^{-\delta d_i'}v_i - \Psi(\theta_i')]\pi(\theta_i') \\ &= \mathbb{E}_{\mathcal{H}_{-i}(\theta_i')})[e^{-\delta d_i'}(v_i' + v_i - v_i') - \Psi(\theta_i')]\pi(\theta_i') \\ &= \mathbb{E}_{\mathcal{H}_{-i}(\theta_i')})[e^{-\delta d_i'}v_i' - \Psi(\theta_i')]\pi(\theta_i') + e^{-\delta d_i'}(v_i - v_i')\pi(\theta_i') \\ &= U(\theta_i') + e^{-\delta d_i'}(v_i - v_i')\pi(\theta_i') \end{aligned} \tag{5.4.18}$$

Since incentive compatibility implies that $U(\theta_i) \geq U(\theta_i'|\theta_i)$, then,

$$U(\theta_i) \geq U(\theta_i') + e^{-\delta d_i'}(v_i - v_i')\pi(\theta_i') \tag{5.4.19}$$

Using (5.4.19) twice, and switching the roles of θ_i and θ_i', we get

$$e^{-\delta d_i'}(v_i - v_i')\pi(\theta_i') \leq U(\theta_i) - U(\theta_i') \leq e^{-\delta d_i}(v_i - v_i')\pi(\theta_i)$$

Letting $\gamma = (v_i - v_i')$, the latter inequality can be rewritten for any $\gamma > 0$ in the following way,

$$e^{-\delta d_i(v_i-\gamma)}\gamma\pi((a_i, v_i-\gamma, d_i(v_i-\gamma))) \leq U((a_i, v_i, d_i))-U((a_i, v_i-\gamma, d_i(v_i-\gamma))) \leq e^{-\delta d_i}\gamma\pi((a_i, v_i, d_i))$$

In virtue of (IC.2), $\pi(\cdot)$ is increasing, then integrable, and therefore

$$U((a_i, v_i, d_i) = U((a_i, \underline{v}, d(\underline{v}))) + \int_{\underline{v}}^{v} e^{-\delta d_i}\pi((a_i, s, d_i(s)))ds$$

For the only if part, notice that (IC.3) plus the fact that $\pi(\cdot)$ is nonnegative, implies (IC). Now, we will show that (5.4.19) follows from (IC.1) and (IC.2).

$$\begin{aligned} U((a_i, v_i, d_i)) &= U((a_i, v_i', d_i')) + \int_{v_i'}^{v_i}(e^{-\delta d_i} - e^{-\delta d_i'})\pi((a_i), s, d_i(s))ds \\ &\geq U((a_i, v_i', d_i')) + (e^{-\delta d_i} - e^{-\delta d_i'})(v_i - v_i')\pi((a_i, v_i', d_i')) \end{aligned}$$

■

5.4.2 Seller's optimality problem

To simplify notation, let ι represent the winner agent in the optimal allocation rule. In other words, ι is such that $y_\iota^{\tau_{w_k}}(s_{w_k}^t) = 1$. Hence, the expected value of the seller at time t, if allocates the object in the mechanism -given a state $s_{w_k}^t$- corresponds to

$$V(s_{w_k}^t|z_{w_k}^t = 0) := \mathbb{E}_{\mathcal{S}_{w_k}|s_{w_k}^t}[\Phi(s_{w_k}|z_{w_k}^t) = 0] + \mathbb{E}_{\mathcal{S}_{w_{k+1}}}[V(s_{w_{k+1}}|z_{w_k}^{\overline{t_k}} = 0)] \tag{5.4.20}$$

where

$$\Phi(s_{w_k}|z^t_{w_k}=0) := e^{-\delta\tau_{w_k}}\left(\alpha_\iota^{\tau_{w_k}}(p_\iota(s_{w_k}^{\tau_{w_k}})-v_0)+(1-\alpha_\iota^{\tau_{w_k}})\left[\int_{a_i}^{\tau_{w_k}}\lambda(d_\iota-t)g(t)\rho^t_\iota dt+V(s^t_{w_k})\right]\right) \tag{5.4.21}$$

The latter expression is the union of two terms. The first one corresponds to seller's expected utility if she allocates the object in the active window w_k given a state $s^t_{w_k}$ (5.4.21), while the second accounts for the continuation value at the beginning of the next window. Specifically, if the seller decides to allocate the object in an arbitrary window w_k, the expected gain is a linear combination of two terms. The first one corresponds to the prices charged to the winning agent, minus seller's reservation value, weighted by the probability α^t_ι that the trade is effectively realized (recall that payments are collected upon allocation). Likewise, the second term corresponds to the fee collected by the seller if the intended-to-being-allocated agent fails to commit with her *ex ante* declared time, plus the continuation value of the seller in the current window (conditioned on the fact that agent ι is unactive). Such event weighted by its correspondent probability $(1-\alpha^t_\iota)$.

Nonetheless, recall that seller can terminate the market at any time to consume the object and get the discounted reservation value, instead to wait for allocating it. Hence, seller will keep market open only if the expected continuation value exceeds her current discounted reservation value (which acts as a lower bound of such valuation).

Therefore, the overall seller's expected values is given by

$$\mathbb{V}(s^t_{w_k}|z^t_{w_k}=0) := \max\{V(s^t_{w_k}|z^t_{w_k}=0), e^{-\delta t}v_0\} \tag{5.4.22}$$

The seller problem consists to choose a direct mechanism Γ that maximizes (5.4.22), given (IC) and (IR).

Proposition 8. *In no optimal mechanism seller terminates the market unilaterally in any dead time.*

Proof. Since no buyer arrives in dead times and no call options are available in the market $V(s^t_{w_k}) = V(s^{\overline{t_k}}_{w_k})$ for all $t \in (\overline{t_k}, \underline{t_k})$. If the seller terminates the market at some $t \in (\overline{t_k}, \underline{t_{k+1}})$, it implies that $e^{-\delta t}v_0 > V(s^{\overline{t_k}}_{w_k})$, but given time is discounted, $e^{\overline{t_k}}v_0 > e^{-\delta t}v_0 > V(s^{\overline{t_k}}_{w_k})$, which directly implies $e^{\overline{t_k}}v_0 > V(s^{\overline{t_k}}_{w_k})$. That is, that seller should close the market at the end of window w_k. ∎

In a nutshell, if the seller decides to terminate the market to consume or execute the good after closing of window w_k, there is no reason to wait later than t_k -the ending bound of the window, given time is discounted. Moreover, if the market is not closed by the seller at time t_k, then she would wait until first buyer arrival in the next window (w_{k+1}) to start making a decision.

Corollary 2. *Given a state s_{w_k}, in window w_k, such that $z_{w_k}^{\overline{t_k}} = 0$ the mechanism arrives to the next window if and only if $\Phi(s_{w_k}|z_{w_k}^{\overline{t_k}} = 0) < \mathbb{E}_{\mathcal{S}_{w_{k+1}}}[V(s_{w_{k+1}}|z_{w_k}^{\overline{t_k}} = 0)]$ and $V(s_{w_k}^{\overline{t_k}}|z_{w_k}^{\overline{t_k}} = 0) > e^{-\delta \overline{t_k}} v_0$.*

Proof. It follows directly from (5.4.20) and proposition 4. ■

Given value is decreasing in time, the expected gain is decreasing along windows, meaning the subset of realizations that would stop the market within the current window is increasing along time. It turns out then, that if seller follow a *threshold rule* -stopping any time the maximum valuation surpasses a bound, then such threshold would be also decreasing with respect to the market window (i.e. buyers become less patient as time goes by).

5.4.3 Implementation of optimal solution

According to Myerson (1981), assumption 1 guarantees that the monotonicity condition (IC.1) is slack at the optimal policy, which means that the allocation rule implemented will correspond to a relaxed solution (see Mierendorff (2010)). Recall that $C_{w_k}(v_i)$ corresponds to buyer i's virtual valuation. Now define, $C_{w_k}^t(v_i) = e^{-\delta t} \alpha_i^t C_{w_k}^t(v_i)$ as the expected virtual valuation for buyer i at any time t. Then, for a given state $s_{w_k}^t$ let $m_{w_k}^t = \max_{i \in A(s_{w_k}^t)} C_{w_k}^t(v_i)$ be the maximal expected virtual valuation among the active buyers in such state. The optimal solution allocates the object to a buyer $i \in \text{argmax}_{i \in A(s_{w_k}^t)} C_{w_k}^t(v_i)$ if the expected virtual valuation of such buyer surpasses the continuation value of reteining the object for the rest of the market window (and for future windows if the object is not allocated in the present window). Formally if $C_{w_k}^t > \mathbb{E}[V(s_{w_k}^t)|z_{w_k}^t = 0]$. In other words, given the object has not been allotted at his arrival and given a history $\mathcal{H}_{-i}$, buyer i (with type θ_i) will be allocated within window w_k if there exists $t^* \in [a_i, d_i^*]$ such that

$$m_{w_k}^t < \mathbb{E}_{\mathcal{H}_{w_k}|\theta_i}[V(s_{w_k}^t)|z_{w_k}^t = 0] \text{ for all } t \in [a_i, t^*] \tag{5.4.23}$$

and

$$C_{w_k}^{t^*}(v_i) = \max\{\mathbb{E}_{\mathcal{H}_{w_k}|\theta_i}[V(s_{w_k}^{t^*})|z_{w_k}^{t^*} = 0], m_{w_k}^{t^*}\} \tag{5.4.24}$$

Equation (5.4.24) says that buyer i's virtual valuation is sufficiently high to induce the seller to allocate him the object at t^*, while (5.4.23) states, precisely, the condition that assures the object is available at such time.

Let $\xi_{a_i}^{d_i^*}(\mathcal{H}_{-i})$ define a *virtual valuation threshold*, given a history $\mathcal{H}_{-i}$, such that if $C_{w_k}^t(v_i) > \xi_{a_i}^{d_i^*}(\mathcal{H}_{-i})$, the buyer is allotted, but not otherwise. Thus,

$$\xi_{a_i}^{d_i^*}(\mathcal{H}_{-i}) = \max\{\ell_{a_i}^{d_i^*}, \max_{j \in A(s_{w_k}^{t^*})\ j \neq i} \{\mathbb{E}_{\mathcal{H}_{w_k}|\theta_i}[V(s_{w_k}^{t^*})|z_{w_k}^{t^*} = 0], C_{w_k}^{t^*}(v_j)\}\} \tag{5.4.25}$$

where $\ell_{a_i}^{d_i^*} = \inf\{x > 0 | m_{w_k}^t < \mathbb{E}_{\mathcal{H}_{w_k}|\theta_i = (a_i, C^{t^{-1}}(x), d_i^*)}[V(s_{w_k}^t)|z_{w_k}^t = 0]\}$ for all $t \in [a_i, t^*]$.

Hence, $\xi_{a_i}^{d_i^*}(\mathcal{H}_{-i})$ is the minimum valuation buyer i has to report to assure the mechanism will arrive at t^*, and to assure he will win at that time. Notice that if it is the case buyer i has the maximum valuation at time t^* he only needs to report a valuation as big as the second higher bidder (as in the static Vickrey auction).

Theorem 2. *Ignoring ties, the optimal solution for the maximization problem of the seller is given by*

$$y_i(s_{w_k}^t) = \begin{cases} 1 & \textit{if} \quad t = t^* \textit{ and } C_{w_k}^t > \xi_{a_i}^{d_i^*}(\mathcal{H}_{-i}) \textit{ for all } t \in [a_i, t^*] \\ 0 & \quad \textit{otherwise} \end{cases} \tag{5.4.26}$$

The payment's rule that supports this allocation pattern as an optimal solution corresponds to

$$p_i(s_{w_k}^t) = \begin{cases} C_{w_k}^{t^{-1}}(\xi_{a_i}^{d_i^*}(\mathcal{H}_{-i})) & \textit{if} \quad z_{w_k}^t = i \\ 0 & \qquad \textit{otherwise} \end{cases} \tag{5.4.27}$$

Proof. The proof follows straightforward. First, with the payment rule described above, payment of a losing buyer is zero, and the winner pays the lowest valuation at which he is able to obtain the object, given a history of reports by the remaining $I - 1$ competitors. Thus, truth-telling is a weakly dominant strategy if buyers follow the optimal report of the deadlines. Hence, the mechanism is that who gives seller the maximum revenue and respects incentive compatibility. ∎

Chapter 6

Application to European Football Transfer Market

The analysis of football markets in Europe has not received the same attention of its counterparts in United States (e.g. american football, baseball, basketball and (ice) hockey). Nonetheless, recently the interest in studying the nature of such market has increased dramatically, spurt in part by the availability of information and the liberalization of labor mobility rigidities. The reasons of why such exploration is a nascent field, correspond mainly to

i) Until recently, individual player salaries were kept secret by teams management, and for long time not even the aggregated wage bill was available to the public

ii) Professional teams have recognized the importance of counting with detailed scientific analysis to make their decisions over when to buy or sell, how much to pay, and what incentives schemes is better for sport's success. At the same time, economists realized that professional team sports offer a unique opportunity to do labor market research, given the degree of detail in which information can be encountered (Frick, 2007).

iii) The unparalleled labor mobility experienced in such market, which exploded after the verdict of the European Court of Justice in the case of Belgian player Jean-Marc Bosman. Such veredict abolished the restriction established in some leagues about the number of foreign-born players allowed to appear in a particular match (Frick, 2007).

Although the interest in studying football markets is relatively new, it has grown at exponential rate in last years, and now many scholars try to disentangle foundational aspects of this

Table 6.1: Transfer expenses and revenues by league 2011-2012

Competition	Expenditure (£ mills)	Revenue (£ mills)
Premier League	562.77	370.55
Serie A	499.95	448.60
Primera División	330.59	277.13
Ligue 1	216.08	162.59
Premier Liga	198.29	60.90
Bundesliga	184.38	152.22

Source: www.transfermarket.co.uk

convulsive market, especially in the relation with transfers dynamics produced each summer in Europe.

As an evidence of the relevance of such market, we present statistics by league on the amounts traded, number of player transferred, and main teams movements. It will be seen that each year there is great variance with respect to the amount of transfers and the time it where realized, which is a clear reflect of the strategic behavior of managers, players and agents.

Using information from www.transfermarket.co.uk correspondent to the season 2011-2012, it can be observed in table 1 that English Premier League and Italian Serie A are at the top of the list with respect to expenditures and revenues (switching positions in each column), followed by Spaniard Primera División, and French Ligue 1.

When disaggregating the information for club in table 2, the top 5 is composed by two English, Chelsea F.C and Manchester City F.C.; one Italian, Juventus F.C, one French, Paris Saint Germaine, and the not common player in the market, Nazi Makhachkala from Russia. The performance of Paris Saint Germain and Manchester City F.C in the market, is a reflection on how the clubs owned by Real Families of Qatar and United Emirates, are bumping the market dramatically. The entering of these new actors (full of liquidity) has made team managers more aggressive when making offers, but at the same time more strategic with respect to what deals accept and at what time.

All these transactions were registered during either the summer or winter transfer window, which are the previously authorized periods by UEFA (Union of European Football Associations) to realize transfers among clubs. Even though *a priori* there are no particular restrictions that differentiates both windows structurally, the majority of transactions are realized in the summer window, due to the necessities of clubs for rearranging the lineups for the incoming season. Additionally, it is worth to remark that summer window lasts for about three months, while

Table 6.2: Top 5 deals, Clubs, 2011-2012

Competition	Expenditure (£ mills)
Paris Saint Germain F.C.	93.50
Chelsea F.C.	90.00
Manchester City F.C.	83.50
Juventus F.C	81.50
Anzhi Makhachkala	74.00

Source: www.transfermarket.co.uk

winter windows last approximately one month. Moreover, there are strategic reasons that make winter window less attractive (since it is at the middle of the season), like the fact that players who have participated with one club in a current Champion or Europe League, are ineligible to play for any other club in the same edition of the competition.
Now, the reasons for which a team can deem necessary go to the market can be very different, but some can be regularly identified such as: qualification to the next Champions or Europe League, retirement of some of its stars, or the announced intention of other mogul club to acquire a key player (which in turn makes necessary go to the market to replace it). All these elements make for managers the process of "buying" or "selling" a player a dynamic game of incomplete information, as presented before.
Actually, in the realm of the model presented here, market windows can be identified with summer market windows (to exclude the asymmetries commented above with respect to winter transfer windows), and the "object" to be allocated is a player[1] who is wanted by I potential clubs. Given that clubs have different stochastic shocks during window's duration (such as injuries or leaves of some players), its entrance in the market is also stochastic, and can be modeled through a counting process $\mathcal{N}$ in each window. As different clubs (potential buyers) have different necessities, they value the abilities of the player in the market, differently, but they learn such value when entering the summer market. All this elements are summarized in valuation v_i, for club i, which is assumed to be an independent draw from F_{w_k}.
The manager of the club who acts as a seller, have the option to place the player in other club and accept the offer they made; or can hold the player for the next season. In such case, they dismiss the opportunity to get the money from the deal, but retain the flow value of player

[1]Notice that player's position is important in the determination of market value, since according to Frick (2007), it is decreasing with the degree of specialization of the position, being goalkeeper the most specialized position, and midfielder the more flexible. However, for the sake of contextualize the theoretical model, we can assume *ex ante* homogeneous players, or we can restrict the attention just to the market of one specific position

services to the team, collapsed in the term v_0. Here, the discount rate over v_0 reflects the fact that as time goes by, the performance of the player is diminishing. Naturally, if the player is not placed within the current window, then the manager has to wait until the opening of the next window.

The seller can observe offer arrivals of any other club, as well as their valuation and the time up to which they are willing to sustain their offers. However, no club can observe which of the rest $I - 1$ competitors is bidding for the player, nor their valuations and reported departure times.

Given that scouting is a normal activity in football markets, managers often receive report about parallel opportunities to get a similar-in-abilities player in not well known markets (Africa, Asia, East Europe or Central America). Those opportunities are received as a signal once the club has entered into the market[2]. When faced with that opportunity, the manager has to decide if takes or dismisses it, under the awareness that if he takes it has to leave the summer window.

In football markets, it is sometimes customary that clubs manifests a deadline up to which they are willing to hold their offers, given that clubs have to do pre-season training, and therefore have to conform their teams as soon as possible. On the light of this fact, many times pre-agreements are signed where the buying club commits to honor the terms signed, otherwise a penalization fee can be collected by the seller. To capture this effect in the present model, a constant penalization fee given by λ is assumed.

With all these elements defined, the manager of a club faces an optimal stopping-problem, as defined in last chapter, whose solution satisfies the results presented above.

[2]Once explanation for this is that information flows more rapidly in Europe than in other emergent markets

Chapter 7

Possible Extensions and Future Research Agenda

The problem analyzed in this paper presents a degree of generality that makes very difficult to address all the tangent questions that arise in one specific paper, yet some of the issues deserve a thorough and deep analysis by themselves. Hence, the purpose of this chapter is to point out some possible extensions and future research agenda the author considers worth to pursue.

Pricing options in dead times. Recall that *inter window* periods are considered dead times at which no buyer can place a bid, nor can they buy call options over a future transaction in the next summer market. Relaxing this assumption supposes the existence of financial markets at which buyers can invest and borrow, as well a legal framework that enforces agents to fulfill all compromises. Moreover, the price of the option would be a not so clear function of all the measures that govern the history of the whole game in the institutionalized market.

Asymmetries in windows. As remarked in the application to European Football Transfer Market, many times market windows are not symmetric, which entails further strategic considerations for sellers and buyers. Those asymmetries go beyond the fact that windows' length are different, but can carry structural differences in valuations, budget restrictions, timing, incompatibilities in the use of the resource, among others. The example of placing a player in the winter market, one he has participated in one of the continental competitions with other clubs, is a real epitome of the strategic intricacies that incorporates the fact to have uneven market windows.

Dynamic budgets. One of the most challenging aspects of dynamic mechanism design is the treatment of budget restrictions. As presented in the review of literature, Athey and Segal (2007) review different theoretical aspects of having balanced budgets with respect to incentive compatibility and efficiency. However, their environment is different from the presented here. Again, inspired in the application of European Market Transfers, many times clubs are in the market for selling and buying, and which event happens first, might affect the other one. Even one a manager enter the market with the only intention to buy, if he happens to receive a good offer for a player, then he could consider selling it and then recompose club's payroll by buying other available players. Clearly, the dynamic interactions it produces should be treated on a more specific environment to maintain tractability. However, its relevance is evident, since on-line markets permit to conduct many transactions in very short periods of time, and electronic methods of payment make possible to use resources almost as soon as operations are realized.

Stochastic Dominance. Observe that in the present model the distribution that governs the i.i.d draws of buyers' valuations can vary along windows, even though the general support of the random variable remains constant. One interesting case of this aspect is considering stochastic dominance of some distribution types. Justification for this assumption can be easily encountered in many environments where the object or entity to be allocated can learn over time and therefore modify their intrinsic value. In our revision of the European Transfer Market, this generalization serves the purpose to model the change in market-value players display, when get more experienced and polish their innate abilities. Actually, many scouts incorporate youngsters to the clubs under the expectation they develop their abilities in a stimulating environment under the guidance and supervision of clubs' coaches. Nonetheless, a natural adaptation time is necessary for players to exploit their whole potential. Once they have passed this time of learning, a time of stability and full productivity is coming, followed by a natural decline in productivity as time goes by. In that sense, players' productivity during their professional-career can be depicted as a bell-shape curve, where the increasing part can be called "the learning stage," the stable portion of the curve at the maximum can be called "mature stage," and the decreasing part can be referred as "decreasing stage." Hence, it is natural that clubs valuations be higher over a player in his "mature stage" than in his "learning" or "decreasing" stage; which can be modeled assuming that distributions (F_{w_k}) in market windows that correspond to the fist case, dominates stochastically those of the second case. Clearly in the introduction of stochastic dominance, would make buyers' and sellers more patient in the "learning stage," and consequently, more aggressive in the "decreasing stage."

All these extensions demand, in some cases, strong modifications of the baseline model, which entails at the same time a deeper exploration on the theory of optimal stopping problems and measure theory. However, its relevance with respect to many online markets make them a worthy endeavor.

Chapter 8

Concluding Remarks

As mentioned before, mechanism design is a branch of theoretical economics that utilizes concepts and inputs from game theory, information economics, dynamic programming, and some key results from real and functional analysis, to analyze the way in which “desirable” social goals can be implemented within democratic and strategic environments. Its origins goes back to the seminal work of Hurwicz (1960, 1972), who set the foundational concepts over which all the subsequent endeavors were established. He introduced the notion of social function and how it can be implemented when agents act selfishly. Later on, in the 1970s and 1980s many advances were made, using the new results coming from the discoveries on games with incomplete information introduced by Harsany (1967, 1968a, 1968b). Specifically, Myerson (1981) and Maskin (1999) contributed to develop the keystone result called “The Revelation Principle,” that permitted to qualify the attention to the class of direct mechanisms. At the same time, theorists continued proving general results with respect to the solution concepts and implementation techiques. The Maskin Monotonicity Theorem, the Graves-Clarke mechanisms for quasilinear environments, and the D’aspermont-Varet result are only examples of how active the field has been since then.

Once established as a solid field within economic theory, new challenges have arised as a consequence of technologies advancements. Probably the most noticeable is the presence of online markets, that permits to a very high number of sellers and buyers meet to trade with very low transaction costs. The epitome of this phenomenon is perhaps the ubiquitous auction sites in the world wide web. The reason of this fact is not arbitrary, since an auction is a better mechanism than over-the counter markets where there is a great variance in the valuation over the object sold, because it enables the seller to extract buyer’s surplus almost entirely. Such characteristic of auction is called in some literature a "discovery effect." However, there is no free lunch, and such discovery effect, generally comes at the price of higher transaction costs, but with the advent of internet schemes these fees practically vanishes and then auctions have become a very popular yet efficient way to trade any kind of commodities, from tickets for a concert to houses and buildings.

As a consequence, much literature has burgeoned to treat with the theoretical issues of on-line auctions, where anonymity offers sellers new tools to increase revenue. Nonetheless, online auctions are not the only dynamic mechanism. Actually many examples can be encountered in public tenders, mechanisms for selling airline tickets, reserving hotels' rooms, and allocate publicity spaces in television. New complexities of problems at hand has obligated the theory to evolve and adopt sophisticated techniques from computer sciences to solve them with heuristic algorithms. Many authors have worked in this applied branch since the explosion of such mechanism, while others have stated in the theoretical side to try to understand several of the new intricacies this dynamic structure imposes over the behavior of agents.

As seen in the review of literature, many papers have developed the general setting of dynamic mechanisms, with special attention to the issues of implementation under incentive compatibility and incentive rationality, either under dominant strategy or Bayesian equilibrium concept. All these endeavors in literature, either modeled in discrete or continuous time, consider complete intervals of time and never analyzes time as a collection of truncated intervals, as it is the case in real-life markets of public tenders, transfer markets of players, or bond auctions at many treasuries or Central Banks. Moreover, another interesting characteristic is that many times deadlines or departure times are not exogenous, and buyers have to decide optimally up to what time they will be willing to stay at the mechanism.

This paper develops a dynamic mechanism with muti-dimmensional space type that allows buyers report an optimal departure time when outside options are available in parallel markets, and where the market is truncated in uniform intervals called as market windows. It extends the work of Mierendorff (2010) who considers exogenous deadlines to analyze the conditions on the relaxed solution is optimal; and the model of Pai and Vohra (2008), who allows complete uncertainty on the side of the seller (i.e. aside of private valuations arrival and departure times are also unknown) to study the implications over incentive compatibility and implementation issues. The main result shows that under general conditions on the optimal stopping rule on the side of the seller, there exists an optimal departure time that balances the trade-off faced by buyers when external trading opportunities are available in parallel markets. Moreover, we demonstrate that if the the seller has the possibility to terminate the market unilaterally, it will never do it in any dead time, but instead will wait to receive new information on the next market window stream. Indeed, at any time t, the opitmal stopping rule always compares the expected revenue starting the next window - conditioned in the history at which the mechanism has arrived- with the potential revenue the seller can raise if allocates the object at such time. Such dicothomy is resolved through a reservation property threshold that depends on the discount rate and on the structure of the space of histories.

Notice that now incentive rationality and incentive compatibility have to take into account the fact that collecting fines is another source of revenue, and hence that buyers internalize the possibility to pay a fine in their revelation problem. However, since the optimal departure time is a function of the private valuation, and thus the uncerainty of the types-space is one-dimmensional, incentive compatibility conditions are an extension of Myerson static version, which is proved.

Finally, implementation of the efficient allocation demands that seller aknowledges the possibility of rent extraction through the collection of fines, versus the effect it produces on the incentive rationality of buyers to participate in the mechanism. Further extension of the model requires more anaylsis on this issue.

Bibliography

[1] Athey, S. and I. Segal, (2007), "An Efficient Dynamic Mechanism," Stanford University, unpublished working paper.

[2] Babioff, M. et al., (2008), "Online auctions and generalized secretary problems," *ACM SIGecom Exchanges*, 7.

[3] Babioff, M. et al., (2007), "A Knapsack Secretary Problem with Applications," in *Proc. 10th Intl. Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, APPROX.

[4] Bergerman, D and J. Valimaki, (2010), "The Dynamic Pivot Mechanism," *Econometrica*, 78, 771-789.

[5] Clarke, E., (1971), "Multipart Pricing of Public Goods," *Public Choice*, 19-33.

[6] Dynkin, E., (1963), "The Optimum Choice of the Instant for Stopping a Markov Process," *Sov. Math. Dokl.* 4.

[7] Frick, B., (2007), "The Football Players' Labor Market: Empirical Evidence from the Major Europeans Leagues," *Scottish Journal of Political Economy*, 54, 422-446.

[8] Gallien, J., (2006), "Dynamic Mechanism Design for Online Commerce," *Operations Research*, 54, 291-310.

[9] Garg, D. et al., (2008), "Foundations of Mechanism Design: A Tutorial Part I- Key Concepts and Classical Results, *Sãdhanã*, 33, 83-130.

[10] Gershkov, A. and B. Moldovanu, (2010), "Incentive Compatible Optimal Stopping," University of Bonn, unpublished working paper.

[11] Gibbard, A. (1973), “Manipulation of Voting Schemes,” *Econometrica*, 41, 587-601.

[12] Groves, T., (1973), “Incentives in Teams,” *Econometrica*, 41, 617-631.

[13] Harsanyi, J. (1967), “Games with Incomplete Information Played by Bayesian Players Part I: The Basic Model,” *Management Sci*, 14, 159-182.

[14] Harsanyi, J. (1968a), “Games with Incomplete Information Played by Bayesian Players Part II: Bayesian Equilibrium Points,” *Management Sci*, 14, 320-334.

[15] Harsanyi, J. (1968a), “Games with Incomplete Information Played by Bayesian Players Part III: The Basic Probability Distribution of the Game,” *Management Sci*, 14, 486-502.

[16] Hurwicz, L., (1960), “Optimality and Informational Efficiency in Resource Allocation Processes,” in *Mathematical Methods in the Social Sciences*, Arrow, K. et al. (editors). Stanford University Press, Palo Alto, USA.

[17] Hurwicz, L., (1972). “On Informationally Decentralized Systems,” in *Decision and Organization*, Radner, R. and C McGuire (editors). North Holland, Amsterdam, Netherlands.

[18] Kleinberg, R., (2005), “A Multiple-Choice Secretary Problem with Applications to Online Auctions,” in *Proc. 16th ACM-SIAM Symposium on Discrete Algorithms*, SODA.

[19] Krishna, V., (2002), *Auction Theory*. Academic Press, New York, USA.

[20] Lange, O., (1936), “On the Economic Theory of Socialism,” *Review of Economic Studies*, 4, 53-71.

[21] Lerner, A., (1944), *The Economics of Control*. McMillan, New York, USA.

[22] Lindley, D., (1961), “Dynamic Programming and Decision Theory. *Applied Statistics*, 10, 3951.

[23] Mas-Colell, A. et al., (1995), *Microeconomic Theory*. Oxford University Press, Oxford, UK.

[24] Maskin, E. and J. Riley, (1984), “Optimal Auctions with Risk Averse Buyers,” *Econometrica*, 52, 1473-1518.

[25] Maskin, E., (1999), “Nash Equilibrium and Welfare Optimality,” *Review of Economic Studies*, 66, 23-38.

[26] Maskin, E., (2008), "Eric Maskin's Nobel Memorial Prize Lecture," Nobel Prize Foundation, Stockholm, Sweeden.

[27] Mierendorff, K., (2011a), "The Dynamic Vicrey Auction," University of Bonn, unpublished working paper.

[28] Mierendorff, K., (2011b), "Optimal Dynamic Mechanism Design with Deadlines," University of Bonn, unpublished working paper.

[29] Milgrom, P. (1989), "Auctions and Bidding: A Primer," *Journal of Economic Perspectives*, 3(3), 3-22.

[30] Myerson, R., (1981), "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58-63.

[31] Myerson, R., (1986), "Multistage Games with Communication," *Econometrica*, 54, 323-358.

[32] Oksendal, B., (1998), *Stochastic Differential Equations: An Introduction with Applications*, fifth edition. Springer-Verlag, Berlin, Germany.

[33] Pai, M. and R. Vohra, (2008), "Optimal Dynamic Auctions," Northwestern University, unpublished working paper.

[34] Pavan, A., et al., (2008), "Dynamic Mechanism Design: Revenue Equivalence, Profit Maximization and Information Disclosure," Northwestern University, unpublished working paper.

[35] Parkes, D. and S. Singh, (2003), "An MDP-Based Approach to Online Mechanism Design," in *Proceedings 17th Annual Conference on Neural Information Processing Systems (NIPS'03).*

[36] Parkes, D., (2007), "Online mechanisms," in *Algorithmic Game Theory*, N. Nissan et al. (editors). Cambridge University Press, Cambridge, UK.

[37] Riley, J. and W. Samuelson (1981), "Optimal Auctions," *American Economic Review*, 71(3), 381-392.

[38] Rochet J. and L. Stole, (2002), "Non-linear Pricing with Random Participation," *Review of Economics Studies*, 69, 277-311.

[39] Siegel, R. (2008), “All Pay Contests,” *Econometrica*, 77, 71-92.

[40] Satterthwaite, M., (1975), “Strategy-Proofness and Arrows Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions,” *Journal of Economic Theory*, 10, 187-217.

[41] von Hayek, F., (1944), *The Road to Serfdom*. Routledge, London, UK.

[42] von Mises, L., (1935), “Die Wirtschaftsrechnung im Sozialistischen Gemeinwesen,” in *Collectivist Economic Planning*, F. von Hayek (editor). Routledge, London, UK.

[43] Zuckerman, D., (1986), “Optimal Stopping in a Continuous Search Model,” *Journal of Applied Probability*, 23, 514-518.