Discussion Paper Series – CRC TR 224

Discussion Paper No. 616
Project B 02

An Efficient Dynamic Mechanism with Covert Information Acquisition

Vitali Gretschko [1]
Jasmina Simon [2]

December 2024

[1] Vitali Gretschko, University of Münster, Email: Vitali.Gretschko@wiwi.uni-muenster.de
[2] Jasmina Simon, University of Mannheim, ZEW, Leibnitz Centre for European Economic Research, Mannheim, Email: Jasmina.Simon@zew.de

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

An Efficient Dynamic Mechanism with Covert Information Acquisition*

Vitali Gretschko† Jasmina Simon‡

12/12/24

Abstract

We examine a setting of independent private value auctions where bidders can covertly acquire gradual information about their valuations. We demonstrate that a dynamic pivot mechanism implements the first-best information acquisition and allocation rule. We apply our results to a commonly used model of auctions with information acquisition. The bidders are symmetric and information acquisition costs are moderate. Our analysis shows that the Dutch auction achieves near-efficiency. That is, the welfare loss is bounded by the information acquisition cost of a single bidder. In contrast, the English auction may result in greater welfare losses.

JEL Classification: D44, D82, D83

Keywords: Information acquisition, dynamic auctions, dynamic pivot mechanism

*We thank seminar participants at the University of Mannheim, Yale, ZEW, and the 12th CRC retreat 2024 for valuable comments. Jasmina Simon acknowledges support by the German Research Foundation (DFG) through CRC TR 224 (Project B02).

†University of Münster. Vitali.Gretschko@wiwi.uni-muenster.de

‡University of Mannheim and ZEW—Leibniz Centre for European Economic Research, Mannheim. Jasmina.Simon@zew.de

1 Introduction

Efficient resource allocation often relies on the ability and willingness of market participants to acquire information. In auction settings, acquiring information is crucial for participants to make informed bidding decisions, which ultimately affects the efficiency of the allocation process. For example, when a government awards contracts for advanced military systems, defense contractors may invest significant resources to evaluate production costs, technological feasibility, and supply chain requirements to estimate their delivery costs. Similarly, in renewable energy support auctions, bidders need to conduct detailed feasibility studies and environmental assessments to accurately determine their costs. This process includes estimating the cost of delivering a renewable energy project, accounting for the opportunity cost of decommissioning an existing power plant, forecasting future revenues from a wind farm, and evaluating potential results of future auctions. Such information acquisition can be costly, involving both direct expenditures and the allocation of significant organizational resources.[1]

Welfare maximization requires careful consideration of the costs associated with information acquisition. It is inefficient for all bidders to invest heavily in information acquisition, as the associated costs are wasted for those who do not win the auction. In this context, it has been established that dynamic standard auctions (English, Dutch) are more efficient than static auctions (first-price, second-price) when information acquisition is considered, as bidders can adjust based on information revealed during the auction. However, whether standard auctions can achieve first-best efficiency remains an open question.

In this paper, we demonstrate in a private value setting that a dynamic first-best mechanism maximizing social welfare exists. Specifically, we consider an environment where a mechanism designer allocates a single indivisible good and maximizes social welfare. Buyers can gradually and covertly acquire private signals to learn about their value for the good. The mechanism designer must balance the cost of information acquisition against the potential benefits of making a more informed allocation decision. To navigate this trade-off, the first-best mechanism, based on the buyers' reports, offers recommendations

[1] According to the International Renewable Energy Agency (2019), by the end of 2018, more than 100 countries had used auctions to purchase renewable energy, representing a ten-fold increase in just one decade.

on information acquisition and determines the allocation of the good accordingly.

To implement the optimal policy, the mechanism needs to satisfy both incentive compatibility and individual rationality constraints. Since information acquisition is both endogenous and unobservable, incentive compatibility constraints are divided into two types: obedience constraints and truth-telling constraints. Obedience constraints ensure that buyers adhere to the mechanism's recommendations for information acquisition, while truth-telling constraints require buyers to truthfully report the signals they acquire.

We prove that the optimal policy can always be implemented using the transfer rule of a dynamic pivot mechanism introduced by Bergemann and Välimäki (2010). As a dynamic version of the Vickrey-Clarke-Groves (VCG) mechanism, a dynamic pivot mechanism requires players to pay the externality they impose on social welfare. As a result, the payment of a bidder does not depend directly on their report except through the allocation decision of the mechanism. This incentivizes buyers to report their signals truthfully. We also demonstrate that the mechanism incentivizes obedience. Specifically, buyers acquire information as recommended, thereby following the optimal information acquisition path. Obedience is crucial to our analysis, as buyers' decisions to acquire information are both endogenous and unobserved.

We use our result to shed light on welfare in standard auctions with information acquisition. Although previous studies (Compte & Jehiel, 2007; Miettinen, 2013; Gretschko & Wambach, 2014; Rezende, 2018) have established that dynamic auctions like the Dutch and English auctions outperform static ones, it remains unclear which of the two is more preferable in terms of welfare. We adopt a framework commonly used in the literature on the dominance of dynamic auctions: buyers' values are symmetrically distributed, and information acquisition is one-shot and fully informative. When the cost of information acquisition is identical for all buyers and sufficiently low, the first-best information acquisition and allocation policy can be described by adapting the results of Doval (2018) on optimal search to an auction setting.

We show that the first-best policy, which is implementable given our initial results, resembles the equilibrium of the Dutch auction studied by Gretschko and Wambach (2014), where buyers acquire information in decreasing order of expected values. We demonstrate that the optimal policy deviates from the Dutch auction equilibrium at most with respect to the optimal information acquisition strategy of the buyer with the lowest ex-ante expected value. Thus, the welfare loss in the Dutch auction, relative to the first-best mechanism, is bounded by the information acquisition cost of a single buyer. By contrast, in the English auction, buyers acquire information in increasing order of expected values, resulting in potentially larger welfare losses.

This result is particularly interesting because the English auction typically allocates efficiently in many settings. In contrast, the Dutch auction is often regarded as less favorable. We show that the Dutch auction outperforms the English auction in our setting with respect to welfare. When bidders are symmetric, the Dutch auction closely resembles the first-best mechanism. This challenges conventional wisdom around auction formats and highlights the potential of the Dutch auction to achieve near-optimal outcomes.

Related literature Our paper contributes to the literature on information acquisition in mechanisms. Several studies examine settings in which buyers can flexibly decide how much and which information to acquire prior to participating in an auction (Bergemann & Välimäki, 2002; Bobkova, 2024; Shi, 2012; Kim & Koh, 2022). Bergemann and Välimäki (2002) also use a Vickrey-Clarke-Groves (VCG) mechanism and show that it ensures efficient information acquisition prior to the mechanism and efficient ex-post allocation of the good. All these papers consider a static environment where information can be acquired *before* the start of the mechanism, once the mechanism starts, information is fixed. These papers also fall into a broader literature on two-step mechanisms, which also encompass models of entry (McAfee & McMillan, 1987; Hausch & Li, 1993; Levin & Smith, 1994; Stegeman, 1996), models of indicative bidding (Ye, 2007; Quint & Hendricks, 2018; Lu, Ye, & Feng, 2021), and models in which the seller controls the disclosure of information (Bergemann & Pesendorfer, 2007; Eső & Szentes, 2007). In contrast, our model allows for flexible information acquisition *during* the mechanism.

Another strand of literature examines dynamic mechanisms in which information evolves exogenously over time (Bergemann & Välimäki, 2010; Athey & Segal, 2013; Pavan, Segal, & Toikka, 2014). A key contribution of our work is to endogenize the flow of information. We build on ideas introduced by Bergemann and Välimäki (2010), who developed a dynamic version of a VCG mechanism known as a dynamic pivot mechanism. In their framework, information evolves over time, and buyers learn their values for free. By contrast, we assume that buyers can choose to acquire information at a cost. The optimal information acquisition path is determined by the mechanism and, unlike in the aforementioned studies, buyers must be incentivized to adhere to this policy. We demonstrate that a dynamic pivot mechanism provides these incentives for obedience, enabling the implementation of the first-best information and allocation policy.

Crémer, Spiegel, and Zheng (2009) is a closely related study that also addresses costly information acquisition during a mechanism. The authors show that an incentive-feasible mechanism exists that implements any first-best policy in their setting. By allowing the seller to charge an entry fee, they obtain a full surplus extraction result. Our model differs in two decisive ways. In Crémer et al. (2009), bidders have no prior information, information is fully informative — buyers need to acquire information only once — whereas we allow for multi-stage and imperfect learning. In their framework, information is also 'productive'; that is, a buyer will never bid without having learned his value, whereas information acquisition is optional in our model.

Our second result examines the performance of standard auction formats relative to the first-best policy in a symmetric setting with moderate information acquisition costs. Several papers show how the design of the auction affects the results and that dynamic auction formats often outperform static formats (Engelbrecht-Wiggans, 1988; Persico, 2000; Klemperer, 2002; Compte & Jehiel, 2004, 2007; Bulow & Klemperer, 2009; Roberts & Sweeting, 2013; Miettinen, 2013; Gretschko & Wambach, 2014; Rezende, 2018; Kleinberg, Waggoner, & Weyl, 2018). Using a mechanism design approach, we demonstrate that the first-best policy in a setting where buyers can covertly acquire information is inherently a *dynamic* mechanism.

In a symmetric setting with moderate information acquisition costs, we derive a first-best information and allocation policy by adapting the optimal search results by Doval (2018) to an auction setting. We are not the first to relate information acquisition in sequential mechanisms to results from optimal search (e.g. Crémer et al., 2009; Kleinberg et al., 2018; Ben-Porath, Dekel, & Lipman, 2024). Kleinberg et al. (2018) examine an auction setting similar to our symmetric setup, applying the optimal search results of Weitzman (1979). In contrast to our assumptions, they require buyers to acquire information before bidding in the auction. In particular, they do not assume that information acquisition is covert. Our results differ in two dimensions: information acquisition is *optional* and *unobservable*. Since information acquisition is optional for auction participation, allocating to a bidder who has not acquired information may be optimal, meaning that the Dutch auction is not always first-best efficient. When information acquisition is unobservable, a mechanism must incentivize buyers to acquire information as recommended. Our first result states that a dynamic pivot mechanism sets the right incentives for buyers to follow the first-best information policy. In a symmetric setting with moderate information costs, Gretschko and Wambach (2014) show that the Dutch auction provides the correct incentives for information acquisition for all buyers but the one with the lowest expected value.

2 Model

Payoffs and priors. There is one seller (she) selling one unit of an indivisible good. There are n buyers (he), $\mathcal{N} = \{1, \ldots, n\}$. The allocation of the good is denoted by $x \in \Delta(\{0,1\}^n)$. That is, x_i is the probability that buyer i receives the good, and thus $\sum_{i=1}^{n} x_i \leq 1$.

A buyer's value ω_i for the good is drawn from a compact set $\Omega_i \subset \mathbb{R}_0^+$. The set of all possible states of the world is $\Omega = \times_{i=1}^{n} \Omega_i$. An element $\omega \in \Omega$ is a vector of values for all buyers, that is $\omega = (\omega_1, \ldots, \omega_n)$. For any vector $(v_1, \ldots, v_n)$ we denote by v_{-i} the vector $(v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n)$ and sometimes write (v_i, v_{-i}) for v. Endow Ω with the Borel σ-Algebra and denote by $\Delta(\Omega)$ the set of all probability distributions on Ω. The prior distribution $\mu(\omega) \in \Delta(\Omega)$ is common knowledge among the seller and the buyers. The marginal distribution over ω_i is denoted by $\mu_i(\omega_i)$. We assume that the prior distribution is independent across buyers, that is

$$\mu(\omega) = \prod_{i=1}^{n} \mu_i(\omega_i).$$

Players do not know anything beyond the prior distribution.

The profit of buyer i depends on the allocation x and a transfer payment to the seller $p_i \in \mathbb{R}$, as follows

$$x_i \omega_i - p_i.$$

For simplicity, the seller's value for the good is 0, so the seller's profit is $\sum_{i=1}^{n} p_i$.

Information acquisition. Each buyer i can acquire information sequentially by receiving up to l_i signals. Signals acquired by different buyers are independent of each other. For $j \in \{1, \ldots, l_i\}$ let S_i^j be a sufficiently rich compact set of possible signal realizations $s_i^j \in S_i^j$. Signal s_i^j is determined by a joint distribution on $\Omega_i \times S_i^1 \ldots S_i^j$. That is, if a buyer chooses to acquire k signals, his belief about ω_i is given by $\pi_{i,k}(\omega_i \mid s_i^1, \ldots, s_i^k)$. Considered as a family of distributions on Ω_i parameterized by s_i^j's, we assume that $\pi_{i,k}(\omega_i \mid s_i^1, \ldots, s_i^k)$ is continuous in each s_i^j in the weak topology on Ω_i.

Acquiring information is costly. Denote by $c_i(k)$ the overall cost of acquiring k signals. The cost difference between acquiring k signals and j signals is given by $\kappa_i(k, j) = c_i(k) - c_i(j)$.

We are interested in sequential information acquisition by all buyers. We define a sequence of information acquisition decisions indexed by $t \in \{0, \ldots, T\}$ as follows. We refer to t as the period. In each period t, any subset of buyers may acquire additional signals. The number of signals acquired by buyer i in all periods prior to and including t is denoted by $\alpha_{i,t}$. The belief of buyer i about ω_i at the beginning of period t depends on all previously acquired signals and we denote it by $\pi_{i,t}(\omega_i \mid s_i^1, \ldots, s_i^{\alpha_{i,t-1}})$. For period t, let $\mathbf{s}_t = (s_1^1, \ldots, s_n^{\alpha_{n,t}})$ represent the profile of all signal realizations for all buyers, and let $\alpha_t = (\alpha_{1,t}, \ldots, \alpha_{n,t})$ denote the profile of all acquisition choices up to and including period t. Let $\pi_t(\omega \mid \mathbf{s}_{t-1}, \alpha_{t-1})$ denote the belief about ω at the beginning of period t. By the independence of the prior belief and the signals, it can be expressed as

$$\pi_t(\omega \mid \mathbf{s}_{t-1}, \alpha_{t-1}) = \prod_{i=1}^{n} \pi_{i,t}(\omega_i \mid s_i^1, \ldots, s_i^{\alpha_{i,t-1}}).$$

Denote the allocation in period t by $\mathbf{x}_t$ and the transfer by $\mathbf{p}_t$. The expected profit of buyer i in period t is given by

$$x_{i,t} \int_{\Omega} \omega_i \, d\pi_{i,t}(\omega_i \mid s_i^1, \ldots, s_i^{\alpha_{i,t-1}}) - c_i(\alpha_{i,t}) - p_{i,t}.$$

Welfare maximization. Agents do not discount the future, and the socially efficient policy maximizes the expected sum of valuations minus the cost of information acquisition. Let $F^{\alpha_t}(\mathbf{s}_t \mid \mathbf{s}_{t-1})$ be the distribution induced on signals in period t by the vector of information acquisition decisions α_t, given the choice of information acquisition α_{t-1} and signal realizations $\mathbf{s}_{t-1}$. The socially optimal program, beginning in period t with belief $\pi_t(\cdot \mid \mathbf{s}_{t-1}, \alpha_{t-1})$ is recursively defined as

$$\begin{aligned} W_t(\mathbf{s}_{t-1}, \alpha_{t-1}) = \max_{x_t, \alpha_t} \sum_{i=1}^{n} \left(x_{i,t} \int_{\Omega} \omega_i \, d\pi_t(\omega \mid \mathbf{s}_{t-1}, \alpha_{t-1}) \right) + \left(1 - \sum_{i=1}^{n} x_{i,t} \right) \\ \times \left(- \sum_{i=1}^{n} \kappa_i(\alpha_{i,t}, \alpha_{i,t-1}) + \int W(\mathbf{s}_t, \alpha_t) \, dF^{\alpha_t}(\mathbf{s}_t \mid \mathbf{s}_{t-1}) \right). \end{aligned} \tag{1}$$

In every period t, the good is either allocated based on the information from period $t-1$ or buyers acquire additional information. Note that $\kappa_i(\alpha_{i,t}, \alpha_{i,t-1}) = c_i(\alpha_{i,t}) - c_i(\alpha_{i,t-1})$ captures the marginal cost of information acquisition in period t for buyer i, which is zero if no information is acquired by this buyer that period. The optimal policy is a tuple $(\mathbf{x}^*, \alpha^*)$, where $\mathbf{x}^* = \{x_t^*\}_{t=0}^T$ represents the first-best allocation, and $\alpha^* = \{\alpha_t^*\}_{t=0}^T$ denotes the corresponding information acquisition policy.

Let T denote the period in which the good is allocated. We restrict our attention to first-best policies where, in every period $t < T$, exactly one agent acquires information. This is without loss as there is no discounting between periods. The number of signal draws is finite for every buyer. Thus, the number of possible information acquisition sequences is finite. As signals are drawn from a compact set, a solution to (1), and thus a first-best policy $(\mathbf{x}^*, \alpha^*)$, always exists. In the optimal policy, the good is allocated to a buyer with the highest expected value for the good given all realized information at time T.

Histories and Mechanism. We focus on direct mechanisms that truthfully implement the socially efficient policy $(\mathbf{x}^*, \alpha^*)$. On a high level, a direct dynamic mechanism works as follows: in each period t, the mechanism asks the buyers to report their most recently acquired signals. Based on the report from period t and the history of reports from all previous periods, the mechanism either allocates the good according to $\mathbf{x}^*$ or, if the good is not allocated, requests the buyers to acquire additional information according to α^*. Potential transfers are realized at the end of each period. Buyers subsequently decide whether to follow the mechanism's recommendation and acquire information before the next period begins.

More formally, the report of buyer i in period t, denoted by r_i^t, reflects the signal he received in period $t-1$. The public history at the start of period t is a tuple comprising all previous reports and information acquisition recommendations $h_t^{\text{pub}} = (r_{t-1}, \alpha_{t-1})$ where $r_0 = \emptyset$ and $r_t = (r_1^1, \ldots, r_n^t)$ for $t > 0$. Let H_t^{pub} denote the set of all possible public histories in period t.

A mechanism is a collection of allocations, information acquisition recommendations, and transfers $\{x_t, \alpha_t, p_t\}_{t=0}^T$. The allocation $x_t : H_t^{\text{pub}} \times S_1^{\alpha_{1,t-1}} \times \ldots \times S_n^{\alpha_{n,t-1}} \times \{\emptyset\} \to \Delta(\{0,1\}^n)$ maps the public history and reports in period t into a probability distribution over allocations. The information acquisition recommendation $\alpha_t : H_t^{\text{pub}} \times S_1^{\alpha_{1,t-1}} \times \ldots \times S_n^{\alpha_{n,t-1}} \times \{\emptyset\} \to \times_{i=1}^n \{\alpha_{i,t-1}, \ldots, l_i\}$ maps the public history and reports in period t into an information acquisition recommendation. The transfer $p_t : H_t^{\text{pub}} \times S_1^{\alpha_{1,t-1}} \times \ldots \times S_n^{\alpha_{n,t-1}} \times \{\emptyset\} \to \mathbb{R}^n$ maps the public history and reports in period t into a vector of transfers for all agents.

We assume pessimistic off-path beliefs. If a buyer reports a signal that is outside his feasible set or refrains from reporting when prompted, the mechanism designer assumes the buyer is of the worst possible type and will never allocate the good to this agent. If a buyer who was not asked to acquire information reports a signal, the mechanism designer disregards the report.

Strategy. A strategy of buyer i in a mechanism $\{x_t^*, \alpha_t^*, p_t\}_{t=0}^T$ in period t is a tuple $(r_i^t, a_{i,t})$, where r_i^t denotes his report in period t, and $a_{i,t}$ indicates how many signals he acquired up to and including period t. The private history of buyer i in period t consists of the public history, the number of signals acquired by the buyer up to period t, and their realizations, so $h_{i,t}^{\text{priv}} = (h_t^{\text{pub}}, a_{i,t-1}, s_i^1, \ldots, s_i^{a_{i,t-1}})$. Let $H_{i,t}^{\text{priv}}$ denote the set of of all possible private histories of buyer i in period t.

A pure reporting strategy for agent i in period t is a mapping from the private history into the relevant signaling space, $r_i^t : H_{i,t}^{\text{priv}} \to S_i^{\alpha_{i,t-1}^*} \cup \{\emptyset\}$. A pure information acquisition strategy is a mapping from the private history into the number of acquired signals, given i has already acquired $a_{i,t-1}$ signals, $a_{i,t} : H_{i,t}^{\text{priv}} \to \{a_{i,t-1}, \ldots, l_i\}$.

At private history $h_{i,t}^{\text{priv}}$, buyer i maximizes his continuation value when deciding on an information acquisition and reporting strategy, as defined below. Let $r_{i,t} = (r_i^1, \ldots r_i^t)$ denote the vector of all signal reports by buyer i up to period t, and let $s_{-i,t} = (s_1^1, \ldots s_{i-1}^{\alpha_{i-1,t}}, s_{i+1}^1, \ldots s_n^{\alpha_{n,t}})$ denote the vector of signals received by all buyers except i up to period

t. Given the mechanism $\{x_t, \alpha_t, p_t\}_{t=0}^{T}$, and truthful reporting by the other buyers, let $F^{a_{i,t}}(\mathbf{s}_t \mid \mathbf{s}_{t-1})$ denote the distribution induced on signals of buyer i in period t by the vector of his information acquisition decision $a_{i,t}$, given the choice of previous information acquisitions $a_{i,t-1}$ and signal realizations $\mathbf{s}_{t-1}$. The optimal reporting strategy $\mathbf{r}_i = \{r_i^t\}_{t=1}^{T}$ and the optimal information acquisition strategy $\mathbf{a}_i = \{a_{i,t}\}_{t=0}^{T}$ for buyer i can be defined recursively.

$$\begin{aligned} V_i(h_{i,t}^{\text{priv}}) = &\max_{r_i^t, a_{i,t}} x_{i,t}(r_{i,t-1}, r_i^t, s_{-i,t-1}, \alpha_{t-1}) \int_{\Omega_i} \omega_i \, d\pi_{i,t-1}(\omega_i \mid h_{i,t}^{\text{priv}}) \\ &- \left(\sum_{i=1}^{n} x_{i,t}(r_{i,t-1}, r_i^t, s_{-i,t-1}, \alpha_{t-1}) \right) p_i(r_{i,t-1}, r_i^t, s_{-i,t-1}, \alpha_{t-1}) \\ &+ \left(1 - \sum_{i=1}^{n} x_{i,t}(r_{i,t-1}, r_i^t, s_{-i,t-1}, \alpha_{t-1}) \right) \Bigg(- \kappa(a_{i,t}, a_{i,t-1}) \\ &+ \int V_i(h_{i,t+1}^{\text{priv}}) \, dF^{a_{i,t}}(\mathbf{s}_t \mid \mathbf{s}_{t-1}) \Bigg). \end{aligned}$$

3 Implementation

We show that the first-best policy $(\mathbf{x}^*, \alpha^*)$ can be implemented. To achieve this, we build on ideas by Bergemann and Välimäki (2010) and employ a variant of the dynamic pivot mechanism to construct a transfer rule that satisfies incentive compatibility and individual rationality conditions. The dynamic pivot mechanism represents a dynamic extension of the Vickrey-Clarke-Groves (VCG) mechanism.

Incentive compatibility encompasses both obedience and truth-telling conditions. The obedience conditions are satisfied if buyers acquire information according to the optimal information acquisition path α^*, meaning they follow the mechanism's recommendations. The truth-telling conditions require that buyers report their signals s_i^j truthfully. Individual rationality requires that a buyer's expected payoff to exceed his outside option, which we normalize to zero.

In a dynamic setting, the concepts of ex post incentive compatibility and individual rationality are defined on periodic basis. A mechanism is *periodic ex post incentive compatible* if, at the time a buyer is asked to acquire information and to report his signal, his best response is to follow the mechanism's recommendation and report truthfully, given that the other buyers do the same. A dynamic mechanism is *periodic ex post individually* *rational* if, at the time of reporting, a buyer's expected payoff exceeds his outside option zero. Note that the report is *ex post* with respect to all the information acquired up to the current period but not necessarily for information that may be acquired in future periods.

The main idea of a static VCG mechanism is that buyers internalize the externality they impose on other buyers. By paying the externality they impose on others, buyers receive a payoff equal to their marginal contribution to social welfare. Consequently, the transfer incentivizes buyer i to maximize social welfare and report his value truthfully. This logic extends to the dynamic version of the mechanism. The main challenge in this setting is to incentivize information acquisition according to the welfare-maximizing policy.

Recall that $W_t(\mathbf{s}_{t-1}, \alpha_{t-1})$ represents social welfare in period t, given the belief $\pi_t(\omega \mid \mathbf{s}_{t-1}, \alpha_{t-1})$, under the optimal allocation and information acquisition policy $(\mathbf{x}^*, \alpha^*)$. Similarly, define the social value excluding buyer i as

$$\begin{aligned} W_{-i,t}(\mathbf{s}_{t-1}, \alpha_{t-1}) = \max_{x_t, \alpha_t} \sum_{j \neq i} \left(x_{j,t} \int_\Omega \omega_j \, d\pi_t(\omega \mid \mathbf{s}_{t-1}, \alpha_{t-1}) \right) + \left(1 - \sum_{j \neq i} x_{j,t} \right) \\ \times \left(- \sum_{j \neq i} \kappa_j(\alpha_{j,t}, \alpha_{j,t-1}) + \int W_{-i}(\mathbf{s}_t, \alpha_t) \, dF^{\alpha_t}(\mathbf{s}_t \mid \mathbf{s}_{t-1}) \right). \end{aligned} \tag{2}$$

The welfare-maximizing allocation without buyer i is denoted by $\mathbf{x}^{*,-i} = \{x_t^{*,-i}\}_{t=0}^T$, and the corresponding information acquisition policy by $\alpha^{*,-i} = \{\alpha_t^{*,-i}\}_{t=0}^T$.

The *marginal contribution* of buyer i is the change in social welfare due to their inclusion in the mechanism, and is defined as

$$M_{i,t}(\mathbf{s}_{t-1}, \alpha_{t-1}) = W_t(\mathbf{s}_{t-1}, \alpha_{t-1}) - W_{-i,t}(\mathbf{s}_{t-1}, \alpha_{t-1}). \tag{3}$$

The marginal contribution of buyer i can be decomposed into *flow* marginal contributions $m_{i,t}(\mathbf{s}_{t-1}, \alpha_{t-1})$, defined recursively as follows:

$$\begin{aligned} m_{i,t}(\mathbf{s}_{t-1}, \alpha_{t-1}) =& M_{i,t}(\mathbf{s}_{t-1}, \alpha_{t-1}) \\ & - \left(1 - \sum_{i=1}^n x_{i,t}(\mathbf{s}_{t-1}, \alpha_{t-1}) \right) \int M_{i,t+1}(\mathbf{s}_t, \alpha_t) \, dF^{\alpha_t}(\mathbf{s}_t \mid \mathbf{s}_{t-1}). \end{aligned} \tag{4}$$

We are now ready to define the transfer payments in the dynamic pivot mechanism.

First, assume without loss of generality that transfers are made when the good is allocated. The transfer payment of a buyer reflects the externality he imposes on other buyers. In other words, the transfer payment equals a buyer's value minus his marginal contribution to social welfare. The transfer payment of buyer i at belief $\pi_t(\omega \mid \mathbf{r}_t, \alpha_{t-1})$ is defined as [2]

$$p_i^*(\mathbf{r}_t, \alpha_{t-1}) = x_{i,t}(\mathbf{r}_t, \alpha_{t-1}) \int_{\Omega} \omega_i \, d\pi_t(\omega \mid \mathbf{r}_t, \alpha_{t-1}) - m_{i,t}(\mathbf{r}_t, \alpha_{t-1}). \tag{5}$$

We are now in the position to state our first result.

Proposition 1. *The transfer rule* $\mathbf{p}^*(\mathbf{r}_t, \alpha_{t-1})$ *implements the first-best information acquisition and allocation rule* $(\mathbf{x}^*, \alpha^*)$*, ensuring periodic ex post incentive compatibility (truth-telling and obedience) and periodic ex post individual rationality.*

Proof. The proof is relegated to the appendix. □

The transfer payment is designed so that every buyer pays the difference between his expected payoff and his flow marginal contribution to social welfare. This key feature of any VCG mechanism ensures that a buyer's transfer does not depend on his report. This incentivizes buyers to report truthfully. In a dynamic setting, we must also address the possibility of multiple joint deviations from truthful reporting. In the proof of Proposition 1, we use a backward induction argument to demonstrate that such deviations are not profitable.

Furthermore, we show that the dynamic pivot mechanism satisfies the obedience conditions, meaning that buyers acquire information as recommended by the mechanism. To see why this is true, observe that the continuation utility of a buyer can be decomposed into flow utility. The flow utility of buyer i in period t is given by:

$$x_{i,t}^*(\mathbf{r}_t, \alpha_{t-1}^*) \int_{\Omega_i} \omega_i \, d\pi_{i,t}(\omega_i \mid s_i^1, \ldots, s_i^{\alpha_{i,t-1}^*}) - \left(\sum_{i=1}^{n} x_{i,t}^*(\mathbf{r}_t, \alpha_{t-1}^*) \right) p_i^*(\mathbf{r}_t, \alpha_{t-1}^*)$$
$$- \left(1 - \sum_{i=1}^{n} x_{i,t}^*(\mathbf{r}_t, \alpha_{t-1}^*) \right) \kappa_i(a_{i,t}, a_{i,t-1}).$$

Note that the flow utility of a buyer matches his flow marginal contribution. Specifically, if the good is not allocated in period t (i.e. $\sum_{i=1}^{n} x_{i,t}^* = 0$), a buyer's flow utility

[2] Note that signals received in period t are reported at the start of the following period, $t+1$; hence, the vector of signals s_{t-1} is reported as r_t.

equals his information acquisition cost, $-\kappa_i(a_{i,t}, a_{i,t-1})$, which also represents his contribution to social welfare. If the good is allocated in period t, a buyer's flow utility equals his expected payoff, $\mathbb{E}[x^*_{i,t}\,\omega_i]$, minus his transfer $p^*_i(\mathbf{r}_{t-1}, \alpha_{t-1})$. Thus, every buyer internalizes his social externality when deciding on information acquisition.

4 The Dutch auction is nearly optimal

Thus far, we havenot explicitly addressed the first-best information and allocation policy. In this section, we apply our insights to a commonly studied setting for analyzing information acquisition in standard auctions. In this setting, bidders possess initial private information about their value and may, at a moderate cost, acquire a fully informative signal. Specifically, acquiring information allows a bidder to learn their value perfectly.

We draw on results of Doval (2018) to derive the first-best information acquisition and allocation strategy for this setting. As shown in Proposition 1, this first-best policy is implementable. We demonstrate that the optimal policy closely resembles the equilibrium strategy of the Dutch auction, as derived in Gretschko and Wambach (2014). In this particular setting, the welfare loss from running a Dutch auction instead of the optimal mechanism is bounded above by the information acquisition cost of a single bidder. Thus, the Dutch auction performs at most slightly worse than the first-best policy in terms of efficiency. By contrast, we will demonstrate that the welfare loss of the English auction can be substantial.

Model. The model consists of n bidders who are imperfectly informed. Each bidder knows his expected value $\mu_i \in [\underline{\mu}, \bar{\mu}]$, which is private information and is drawn from a continuous distribution function G.[3] A bidder's value, ω_i, is defined as the sum of his expected value μ_i and a noise term ϵ_i, such that $\omega_i = \mu_i + \epsilon_i$. We assume that for every bidder the noise parameter ϵ_i is drawn from the same symmetric distribution function F, $\epsilon_i \sim F\ \forall i$. The distribution is symmetric around zero and admits density function f. Thus, the value ω_i of a bidder is distributed according to F, shifted by the prior μ_i, and we write $\omega_i \sim F_i$.

[3] We slightly abuse notation, as in the previous sections μ_i is common knowledge. However, we assume without loss of generality that the first signal is free, and μ_i represents the updated expected value after observing its realization. In this case, the first-best mechanism requires that all bidders acquire this signal, and they will comply since it is free.

Every bidder can acquire one signal and this signal is fully informative. We assume the cost of acquiring this signal is identical for all bidders and we denote it by c. When a bidder decides to acquire information, he learns the realization of the noise parameter ϵ_i, and thus, his value ω_i. Furthermore, we assume that the information acquisition cost is sufficiently small.

Assumption 1. *Assume that*

$$\int_{\bar{\mu}}^{\infty} (\omega_i - \bar{\mu}) dF(\omega_i - \underline{\mu}) > c.$$

Assumption 1 ensures that information acquisition plays a significant role under the first-best policy. Specifically,it implies that the object is almost never allocated to a bidder who has not acquired information. This condition is satisfied if either c is sufficiently small or the uncertainty about the true valuation is sufficiently pronounced as compared to the range of expected values.

First-best mechanism. The first-best policy is derived from Doval (2018). The first-best mechanism implementing this policy first requires bidders to report their expected values μ_i. The bidder with the highest expected value is asked to acquire information first and report his valuation to the mechanism. The first-best information acquisition policy follows a cutoff rule. Specifically, the good is allocated to the bidder with the highest valuation if his valuation exceeds a cutoff determined by the reservation values (defined below) of those who have not acquired information. If only one bidder remains without acquired information and his back-up value (defined below) exceeds the highest reported valuation of the other bidders, the good is allocated to this bidder without requiring additional information acquisition.

The cutoffs for each bidder, as described above, are determined by their *reservation value* ω_i^R and *backup value* ω_i^B. The reservation value, ω_i^R, is defined as

$$c = \int_{\omega_i^R}^{\infty} (\omega_i - \omega_i^R) \, dF_i(\omega_i).$$

Similarly, the backup value ω_i^B, is defined as

$$c = \int_0^{\omega_i^B} (\omega_i^B - \omega_i)\, dF_i(\omega_i).$$

The optimal policy can be formally stated as follows.

Proposition 2 (Theorem 1 in Doval, 2018)**.** *Assume that Assumption 1 holds true. Order the n bidders in decreasing expected value order, such that $\mu_1 \geq \ldots \geq \mu_n$.*

The optimal policy is as follows:

Order *Have bidders acquire information in decreasing expected value order.*

Stopping

1. *Stop if the highest sampled valuation exceeds all non-sampled reservation values, and allocate the good to the bidder with this highest value.*
2. *If only one bidder remains who has not acquired information, stop if the highest sampled valuation is less than his backup value, and allocate the good to the remaining bidder.*

To maximize welfare, if any bidder acquires information it should be the bidder with the highest expected value. This follows from the symmetry of the noise parameter and equal information acquisition costs. Assumption 1 ensures that the reservation value of a bidder with the the lowest possible expected value exceeds the highest possible expected value. In this case, it is never optimal to allocate the good to any bidder other than the n^{th} without first requiring them to acquire information.

The cutoffs define the optimal stopping conditions. At the cutoffs, the mechanism is indifferent between asking a bidder acquire information or choosing the optimal outside option. If the outside option exceeds the bidder's reservation value, ω_i^R, the mechanism will choose the outside option. Specifically, if the value of a bidder who has acquired information exceeds the reservation values of all bidders without information, the mechanism stops and allocates the good. This is reflected in stopping condition 1. When the outside option lies between the cutoffs, the mechanism asks bidder i to acquire information, as the outside option is sufficiently close to bidder i's expected value μ_i. One outside option is the continuation value obtained by having at least one more bidder acquire information. The restriction on the cost parameter ensures that the continuation value prevents allocation of the good to any bidder (except bidder n) who has not acquired information. If the outside option is lower than bidder i's backup value, ω_i^B, the mechanism maximizes social welfare by allocating the good to bidder i without additional information acquisition. This corresponds to stopping condition 2. The assumption on the cost parameter ensures that the continuation value is sufficiently high to prevent allocation of the good to any bidder other than the one with the lowest expected value without acquiring information.

Equilibrium in the Dutch auction In a descending auction, a bidder's strategy involves deciding when to acquire information, if at all, and when to stop the clock, given that any bidder's acceptance ends the auction as the clock price decreases. Gretschko and Wambach (2014) demonstrate that an equilibrium with information acquisition exists in the Dutch auction, and that both equilibrium bidding and information acquisition can be characterized by a single increasing function $\beta(\cdot)$.

Proposition 3 (Gretschko & Wambach, 2014)**.** *There exists $\bar{c}^D > 0$ such that, for all $c \leq \bar{c}^D$, an equilibrium $\beta(\cdot)$ of the descending auction exists with the following properties:*

(i) Bidder i acquires information if and only if the clock price reaches $\beta(\omega_i^R)$.

(ii) If, at the clock price of information acquisition p, the bidder learns that $\beta(\omega_i) \geq p$, he stops the clock immediately;

(iii) If, at the clock price of information acquisition p, the bidder learns that $\beta(\omega_i) < p$, he stops the clock at $\beta(\omega_i)$.

Bidders acquire information in decreasing reservation value order. Due to the symmetry of the noise parameter, the ordering of the reservation values coincides with the ordering of expected values. Therefore, the order of information acquisition is identical in the optimal policy and the Dutch auction. A bidder stops the clock when the clock price equals the value of the bidding function of his value or when he learns that the value of the bidding function of his value exceeds the clock price. Notably, since the bidding function is identical for all bidders, whenever a bidder stops the clock, his value exceeds the reservation values of all bidders who have not yet acquired information. Thus, the stopping condition of the Dutch auction aligns with the first-best stopping condition for all bidders except the one with the lowest expected value. In the Dutch auction, if no bidder stops the clock earlier, the final bidder also acquires information. In the optimal mechanism, the good may be allocated to the bidder with the lowest prior without requiring him to acquire information. Consequently, the expected welfare loss of the Dutch auction relative to the first-best policy is bounded above by the information acquisition cost of a single bidder.

Corollary 1. *For all information acquisition costs $c \leq \bar{c}^D$, the welfare loss of the Dutch auction relative to the first-best mechanism is bounded above by c.*

Overall, if the information acquisition cost is sufficiently small and the bidders ex-ante symmetric, the Dutch auction nearly implements the first-best policy.

Equilibrium in the English auction In an English auction, a bidder's strategy involves determining whether and when to acquire information and deciding when to exit the auction without obtaining the good. Rezende (2018) demonstrates the existence of an equilibrium with information acquisition in the English auction. We focus on sufficiently low information acquisition costs, ensuring that all bidders acquire information provided at least one other bidder is also bidding at the optimal time for information acquisition. In equilibrium, bidders bid up to their respective value. Bidders must decide when to acquire information.

Proposition 4 (Rezende, 2018)**.** *There exists $\bar{c}^E > 0$ such that for all $c \leq \bar{c}^E$ an equilibrium $\beta(\cdot)$ of the English auction exists with the following properties:*

(i) Bidder i acquires information if and only if the clock price reaches ω_i^B.

(ii) If at the clock price of information acquisition p the bidder learns that $\omega_i \geq p$, he will remain in the auction.

(iii) If at the clock price of information acquisition p the bidder learns that $\omega_i < p$, he will drop out of the auction.

When bidders decide to acquire information, they face a trade-off between acquiring it early and delaying it. The benefit of delaying information acquisition lies in potential cost savings, which occur only if the auction concludes during the delay. The risk of delaying information acquisition is that the true value of the item is below the price paid. Bidders acquire information in order of increasing backup values. In this symmetric setting, this implies that all bidders acquire information in order of increasing expected values. The only bidder who may not acquire information is the one with the highest expected value, as he will win without acquiring information if all other bidders drop out before it becomes optimal for him to do so.

Comparing the English auction to the first-best mechanism shows that the welfare loss from the English auction can be as large as $(n-1)c$. This occurs when, in the first-best mechanism, only the bidder with the highest expected value acquires information and discovers a high value, leading to an immediate allocation of the object. In this scenario, the English auction allocates the object to the same bidder. However, before the allocation, all other bidders - potentially including the winner - must acquire information, resulting in the described welfare loss.

Corollary 2. *The welfare loss of the English auction can be as large as* $(n-1)c$*.*

5 Conclusion and discussion

This paper challenges conventional wisdom about auction design, particularly regarding the relative merits of Dutch and English auctions. While English auctions are often favored for their transparency and ability to achieve efficient allocations, our analysis reveals that Dutch auctions can outperform English auctions when bidders face significant information acquisition costs.

The key insight is that the Dutch auction's descending price format naturally implements an information acquisition order that closely mirrors the socially optimal sequence, where bidders with higher expected values acquire information first. This alignment results in a welfare loss that is bounded by a single bidder's information acquisition cost. In contrast, the English auction's ascending format induces bidders to acquire information starting from those with lower expected values, potentially causing substantial inefficiencies as multiple bidders may incur unnecessary information costs before the auction concludes. However, this insight comes with important caveats. The superiority of the Dutch auction relies on the assumptions that bidders are ex-ante symmetric and that information acquisition costs are moderate. In settings where these assumptions fail to hold, different auction formats may be optimal.

The broader takeaway is that coordinating information acquisition can be as important for efficiency as the allocation rule itself. This suggests that mechanism designers should pay careful attention to how different auction formats shape bidders' incentives to acquire information, rather than focusing exclusively on allocation outcomes. For practitioners, our results provide a novel argument for considering Dutch auctions in settings where English auctions have traditionally been preferred.

Appendix

Proof of Proposition 1

We demonstrate that the mechanism $\{x_t^*, \alpha_t^*, p^*\}_{t=0}^T$ satisfies periodic ex post incentive compatibility and individual rationality. For simplicity, we omit the stars in the notation. The proof is structures into three parts. First, we demonstrate that one-shot deviations from truthful reporting and obedience do not yield a profit. Next, we examine multi-step deviations and demonstrate that no profitable deviations exist.

Truthful reporting

Recall that the mechanism considers only those signals that are on path, meaning those that would be acquired according to the optimal information acquisition path α. We prove that agent i is incentivized to report his signal, $s_i^{\alpha_{i,t-1}}$, truthfully in period t. Writing signals as reports whenever buyers tell the truth, the vector $(s_{i,t-2}, r_i^t, s_{-i,t-1})$ denotes the reports by all players until and including period t where all buyers tell the truth in every period except for player i who reports r_i^t in period t. Note that a report in period t reflects a signal received in period $t-1$. Therefore, under truthful reporting $r_i^t = s_i^{\alpha_{i,t-1}}$. Future information acquisition recommendations $\alpha_t, \alpha_{t+1}, \dots$ depend on prior signal reports $(s_{i,t-2}, r_i^t, s_{-i,t-1})$. When a report is not truthful, the recommendation may differ from the recommendation given after a truthful signal. To simplify notation, we denote the possible change in recommendation as $\hat{\alpha}_t$. It suffices to show that a buyer receives his marginal contribution as his continuation value to ensure truthful reporting.

That is, for all $s_i^{\alpha_{i,t-1}} \in S_i^{\alpha_{i,t-1}}$ the following must hold:

$$
\begin{aligned}
& x_{i,t}(s_{t-1},\alpha_{t-1}) \int_{\Omega_i} \omega_i \, d\pi_{i,t}(\omega_i \mid s_i^1,\ldots,s_i^{\alpha_{i,t-1}}) - \left(\sum_{i=1}^n x_{i,t}(s_{t-1},\alpha_{t-1})\right) p_i(s_{t-1},\alpha_{t-1}) \\
& \quad + \left(1 - \sum_{i=1}^n x_{i,t}(s_{t-1},\alpha_{t-1})\right) \Bigg(- \kappa_i(\alpha_{i,t},\alpha_{i,t-1}) \\
& \quad + \int_\Omega M_{i,t+1}(s_t,\alpha_t) \, dF^{\alpha_t}(\mathbf{s}_t \mid \mathbf{s}_{t-1}) \Bigg) \\
& \geq x_{i,t}(s_{i,t-2}, r_i^t, s_{-i,t-1}, \alpha_{t-1}) \int_{\Omega_i} \omega_i \, d\pi_{i,t}(\omega_i \mid s_i^1,\ldots,s_i^{\alpha_{i,t-1}}) \qquad (6) \\
& \quad - \left(\sum_{i=1}^n x_{i,t}(s_{i,t-2}, r_i^t, s_{-i,t-1}, \alpha_{t-1})\right) p_i(s_{i,t-2}, r_i^t, s_{-i,t-1}, \alpha_{t-1}) \\
& \quad + \left(1 - \sum_{i=1}^n x_{i,t}(s_{i,t-2}, r_i^t, s_{-i,t-1}, \alpha_{t-1})\right) \Bigg(- \kappa_i(\hat{\alpha}_{i,t},\alpha_{i,t-1}) \\
& \quad + \int_\Omega M_{i,t+1}(s_{i,t-2}, r_i^t, s_i^{\hat{\alpha}_{i,t}}, s_{-i,t}, \hat{\alpha}_t) \, dF^{\hat{\alpha}_t}(s_t \mid s_{t-1}) \Bigg).
\end{aligned}
$$

The optimal allocation decisions, x_t, and information acquisition recommendations, α_t, may differ depending on whether buyer i participates. Let x_t^{-i} denote the efficient allocations without buyer i, and let α_t^{-i} denote the corresponding efficient information acquisitions.

By construction of $p_i(r_{t-1},\alpha_{t-1})$ in (5), the left hand side of (6) equals buyer i's marginal contribution to social welfare. We express marginal contributions in terms of social values, as defined in (3), and plug in the payment rule $p_i(r_{t-1},\alpha_{t-1})$ to rewrite the right hand side.

$$
\begin{aligned}
&W_t(\mathbf{s}_{t-1}, \alpha_{t-1}) - W_{-i,t}(\mathbf{s}_{t-1}, \alpha_{t-1}) \\
&\geq x_{i,t}(s_{i,t-2}, r_i^t, s_{-i,t-1}, \alpha_{t-1}) \int_{\Omega_i} \omega_i \, d\pi_{i,t}(\omega_i \mid s_i^1, \ldots, s_i^{\alpha_{i,t-1}}) \\
&\quad - \Bigg[- \sum_{j \neq i} \left(x_{j,t}(s_{i,t-2}, r_i^t, s_{-i,t-1}, \alpha_{t-1}) \int_{\Omega_j} \omega_j \, d\pi_{j,t}(\omega_j \mid s_j^1, \ldots, s_j^{\alpha_{j,t-1}}) \right) \\
&\qquad + \left(1 - \sum_{i=1}^{n} x_{i,t}(s_{i,t-2}, r_i^t, s_{-i,t-1}, \alpha_{t-1}) \right) \left(\sum_{j \neq i} \kappa_j(\hat{\alpha}_{j,t}, \alpha_{j,t-1}) \right) \\
&\qquad + \sum_{j \neq i} \left(x_{j,t}^{-i}(s_{i,t-2}, r_i^t, s_{-i,t-1}, \alpha_{t-1}) \int_{\Omega_j} \omega_j \, d\pi_{j,t}(\omega_j \mid s_j^1, \ldots, s_j^{\alpha_{j,t-1}}) \right) \\
&\qquad + \left(1 - \sum_{j \neq i} x_{j,t}^{-i}(s_{i,t-2}, r_i^t, s_{-i,t-1}, \alpha_{t-1}) \right) \left(- \sum_{j \neq i} \kappa_j(\hat{\alpha}_{j,t}^{-i}, \alpha_{j,t-1}) \right) \\
&\qquad - \left(1 - \sum_{j \neq i} x_{j,t}^{-i}(s_{i,t-2}, r_i^t, s_{-i,t-1}, \alpha_{t-1}) \right) \\
&\qquad\quad \times \int_{\Omega} W_{-i,t+1}(s_{i,t-2}, r_i^t, s_i^t, s_{-i,t}, \hat{\alpha}_t^{-i}) \, dF^{\hat{\alpha}_t^{-i}}(s_t \mid s_{t-1}) \\
&\qquad + \left(1 - \sum_{i=1}^{n} x_{i,t}(s_{i,t-2}, r_i^t, s_{-i,t-1}, \alpha_{t-1}) \right) \\
&\qquad\quad \times \int_{\Omega} W_{-i,t+1}(s_{i,t-2}, r_i^t, s_i^t, s_{-i,t}, \hat{\alpha}_t) \, dF^{\hat{\alpha}_t}(s_t \mid s_{t-1}) \Bigg] \\
&\quad + \left(1 - \sum_{i=1}^{n} x_{i,t}(s_{i,t-2}, r_i^t, s_{-i,t-1}, \alpha_{t-1}) \right) \Bigg(- \kappa_i(\hat{\alpha}_{i,t}, \alpha_{i,t-1}) \\
&\quad + \int_{\Omega} W_{t+1}(s_{i,t-2}, r_i^t, s_i^t, s_{-i,t}, \hat{\alpha}_t) \, dF^{\hat{\alpha}_t}(s_t \mid s_{t-1}) \\
&\quad - \int_{\Omega} W_{-i,t+1}(s_{i,t-2}, r_i^t, s_i^t, s_{-i,t}, \hat{\alpha}_t) \, dF^{\hat{\alpha}_t}(s_t \mid s_{t-1}) \Bigg).
\end{aligned}
$$

We rewrite the inequality.

$$
\begin{aligned}
&W_t(\mathbf{s}_{t-1}, \alpha_{t-1}) - W_{-i,t}(\mathbf{s}_{t-1}, \alpha_{t-1}) \\
&\geq \sum_{j=1}^{n} \left(x_{j,t}(s_{i,t-2}, r_i^t, s_{-i,t-1}, \alpha_{t-1}) \int_{\Omega_j} \omega_j \, d\pi_{j,t}(\omega_j \mid s_j^1, \ldots, s_j^{\alpha_{j,t-1}}) \right) \\
&\quad + \left(1 - \sum_{j=1}^{n} x_{j,t}(s_{i,t-2}, r_i^t, s_{-i,t-1}, \alpha_{t-1})\right) \left(- \sum_{j=1}^{n} \kappa_j(\hat{\alpha}_{j,t}, \alpha_{j,t-1}) \right. \\
&\quad \left. + \int_{\Omega} W_{t+1}(s_{i,t-1}, r_i^t, s_{-i,t}, \hat{\alpha}_t) \, dF^{\hat{\alpha}_t}(s_t \mid s_{t-1}) \right) \\
&\quad - W_{-i,t}(\mathbf{s}_{t-1}, \alpha_{t-1}).
\end{aligned}
$$

The above inequality holds for all r_i^t by the social optimality of x_t and α_t.

Obedience

We demonstrate that acquiring information is a best response for a buyer i in period t when the mechanism recommends it. Recall that, in the optimal mechanism, only one buyer is asked to acquire information in any period t, and assume that this buyer is i in period t. Let $\hat{M}_{i,t+1}(s_{t-1}, r_i^{t+1}, \alpha_t)$ denote agent i's marginal contribution starting in period $t+1$, assuming he did not acquire information in period t and reports r_i^{t+1}. Since agent i does not receive a signal in period t, he can report any signal or choose not to report. To incentivize information acquisition in period t, the following condition must hold.

$$
\begin{aligned}
&x_{i,t}(s_{t-1}, \alpha_{t-1}) \int_{\Omega_i} \omega_i \, d\pi_{i,t}(\omega_i \mid s_i^1, \ldots, s_i^{\alpha_{i,t-1}}) - \left(\sum_{i=1}^{n} x_{i,t}(s_{t-1}, \alpha_{t-1}) \right) p_i(s_{t-1}, \alpha_{t-1}) \\
&\quad + \left(1 - \sum_{i=1}^{n} x_{i,t}(s_{t-1}, \alpha_{t-1})\right) \left(- \kappa_i(\alpha_{i,t}, \alpha_{i,t-1}) + \int_{\Omega} M_{i,t+1}(s_t, \alpha_t) \, dF^{\alpha_t}(\mathbf{s}_t \mid \mathbf{s}_{t-1}) \right) \\
&\geq x_{i,t}(s_{t-1}, \alpha_{t-1}) \int_{\Omega_i} \omega_i \, d\pi_{i,t}(\omega_i \mid s_i^1, \ldots, s_i^{\alpha_{i,t-1}}) - \left(\sum_{i=1}^{n} x_{i,t}(s_{t-1}, \alpha_{t-1}) \right) p_i(s_{t-1}, \alpha_{t-1}) \\
&\quad + \left(1 - \sum_{i=1}^{n} x_{i,t}(s_{t-1}, \alpha_{t-1})\right) \left(\hat{M}_{i,t+1}(s_{t-1}, r_i^{t+1}, \alpha_t) \right)
\end{aligned}
$$

As before, we rewrite the equation by plugging in $p_i(s_{t-1}, \alpha_{t-1})$ from (5) and combining terms.

$$\begin{aligned}
&W_t(\mathbf{s}_{t-1}, \alpha_{t-1}) - W_{-i,t}(\mathbf{s}_{t-1}, \alpha_{t-1}) \\
&\geq \sum_{j=1}^{n} \left(x_{j,t}(s_{t-1}, \alpha_{t-1}) \int_{\Omega_j} \omega_j \, d\pi_{j,t}(\omega_j \mid s_j^1, \dots, s_j^{\alpha_{j,t-1}}) \right) \\
&\quad + \left(1 - \sum_{j=1}^{n} x_{j,t}(s_{t-1}, \alpha_{t-1})\right) \Bigg(- \sum_{j \neq i} \kappa_j(\alpha_{j,t}, \alpha_{j,t-1}) \\
&\quad + W_{t+1}(s_{t-1}, r_i^{t+1}, \alpha_t) \Bigg) - W_{-i,t}(\mathbf{s}_{t-1}, \alpha_{t-1}).
\end{aligned}$$

Since buyer i is asked to acquire information in period t, the good is not allocated in this period; therefore $\sum_{i=1}^{n} x_{i,t}(s_{t-1}, \alpha_{t-1}) = 0$. As buyer i is asked to acquire information in period t, no other agent acquires information in this period; therefore $\sum_{j \neq i} \kappa_j(\alpha_{j,t}, \alpha_{j,t-1}) = 0$. Next, we focus on the term $W_{t+1}(s_{t-1}, r_i^{t+1}, \alpha_t)$. The signal space is sufficiently rich such that an uninformative signal s_i^0 exists for every buyer i in any set S_i^j, which does not alter the belief held by the mechanism when reported: $\pi_t(s_{t-1}, \alpha_{t-1}) = \pi_{t+1}(s_{t-1}, r_i^{t+1}, \alpha_t)$. When buyer i does not acquire information, his belief $\pi_{i,t}(s_{t-1}, \alpha_{t-1})$ remains unchanged. Under truthful reporting, the best response is to report the uninformative signal s_i^0 when his information remains unchanged. We now rewrite the inequality.

$$\begin{aligned}
&- \kappa_i(\alpha_{i,t}, \alpha_{i,t-1}) + \int_{\Omega} W_{t+1}(s_t, \alpha_t) \, dF^{\alpha_t}(s_t \mid s_{t-1}) \\
&\geq W_t(s_{t-1}, \alpha_{t-1}) \geq W_{t+1}(s_{t-1}, s_i^0, \alpha_t)
\end{aligned}$$

The first inequality holds by the optimal information acquisition path, α^*. When it is optimal to ask buyer i to acquire information in period t, the expected value doing so and continuing optimally from period $t+1$ onward exceeds the value in period t. The second inequality holds because the belief about the values remains unchanged, $\pi_t(\omega \mid s_{t-1}, \alpha_{t-1} = \pi_{t+1}(\omega \mid s_{t-1}, s_i^0, \alpha_t)$, but the value of search is weakly higher on the left-hand side due to the option to acquire one additional signal.

Multi-step deviations

Examining one-shot deviations is not sufficient. It is also necessary to rule out cases where buyers deviate from the recommended information acquisition or fail to report truthfully over multiple consecutive rounds. Recall that each buyer can acquire a maximum of l_i signals. No buyer will acquire information or report a signal unless recommended to do so, as a result of pessimistic off-path beliefs.

We first consider multi-step deviations. Assume, for the sake of contradiction, that a profitable multi-step deviation exists for player i. The supposed profitable deviation for player i is to misreport his signals p consecutive times when asked to report them. We showed that a one-shot deviation from truthful reporting is not profitable. Thus, the final deviation in the multi-step deviation is not profitable, and the multi-step deviation must yield a weakly larger payoff due to the first $p-1$ deviations. By repeatedly applying the same reasoning, the first deviation in the p deviations must be profitable, which contradicts the earlier result that no one-shot deviations are profitable.

Similarly, we show that no multi-step deviation in acquiring information can be profitable. Assume there exists a multi-step deviation from the optimal information path, where buyer i does not acquire information $p \leq l_i$ times despite being recommended to do so. Suppose we are in the final period of the profitable multi-step deviation, denoted by P. Let $\tilde{M}_{i,P}(\tilde{r}_{i,P}, \tilde{s}_{-i,P-1}, \tilde{\alpha}_{P-1})$ represent the marginal contribution when buyer i has deviated $p-1$ times from the recommendation and reports optimally as $\tilde{r}_P$. We use $\tilde{s}_{-i,P-1}$ and $\tilde{\alpha}_{P-1}$ to emphasize that changes in buyer i's reports may alter the information acquisition recommendations, α_t, and consequently the vector of signals s_t. Since the mechanism designer is unaware of the previous deviations and recommends information acquisition based on her belief, the following holds

$$
\begin{aligned}
&- \kappa_i(\tilde{\alpha}_{i,P}, \tilde{\alpha}_{i,P-1}) + \int_{\Omega_i} \tilde{M}_{i,P+1}(\tilde{r}_{i,P}, \tilde{s}_i^P, \tilde{s}_{-i,P}, \tilde{\alpha}_P)\, dF^{\tilde{\alpha}_P}(\tilde{s}_P \mid \tilde{s}_{P-1}) \\
&\geq \tilde{M}_{i,P+1}(\tilde{r}_{i,P+1}, \tilde{s}_{-i,P}, \tilde{\alpha}_P).
\end{aligned}
$$

Specifically, the expected marginal contribution from acquiring the recommended signal in period P and truthfully reporting it in period $P+1$ exceeds the marginal contribution from not acquiring information and reporting optimally. Thus, the final step of the multi-step deviation is not profitable, implying the previous $p - 1$ deviations must be. By repeating this argument, we establish that for the multi-step deviation to be profitable, the first deviation must be profitable. Since we previously demonstrated that one-shot deviations are not profitable, it follows that no multi-step deviation can be profitable. This concludes the proof.

References

Athey, S., & Segal, I. (2013). An efficient dynamic mechanism. *Econometrica*, *81*(6), 2463–2485.

Ben-Porath, E., Dekel, E., & Lipman, B. (2024). *Sequential mechanisms for evidence acquisition* [Working Paper]. Retrieved from `https://www.bartlipman.org/working-papers`

Bergemann, D., & Pesendorfer, M. (2007). Information structures in optimal auctions. *Journal of Economic Theory*, *137*(1), 580–609.

Bergemann, D., & Välimäki, J. (2002). Information acquisition and efficient mechanism design. *Econometrica*, *70*(3), 1007–1033.

Bergemann, D., & Välimäki, J. (2010). The dynamic pivot mechanism. *Econometrica*, *78*(2), 771–789.

Bobkova, N. (2024). Information choice in auctions. *American Economic Review*, *114*(7), 1883–1915.

Bulow, J., & Klemperer, P. (2009). Why do sellers (usually) prefer auctions? *American Economic Review*, *99*(4), 1544–1575.

Compte, O., & Jehiel, P. (2004). The wait-and-see option in ascending price auctions. *Journal of the European Economic Association*, *2*(2-3), 494–503.

Compte, O., & Jehiel, P. (2007). Auctions and information acquisition: sealed bid or dynamic formats? *The RAND Journal of Economics*, *38*(2), 355–372.

Crémer, J., Spiegel, Y., & Zheng, C. Z. (2009). Auctions with costly information acquisition. *Economic Theory*, *38*, 41–72.

Doval, L. (2018). Whether or not to open Pandora's box. *Journal of Economic Theory*, *175*, 127–158.

Engelbrecht-Wiggans, R. (1988). On a possible benefit to bid takers from using multi-stage auctions. *Management Science*, *34*(9), 1109–1120.

Eső, P., & Szentes, B. (2007). Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies*, *74*(3), 705–731.

Gretschko, V., & Wambach, A. (2014). Information acquisition during a descending auction. *Economic Theory*, *55*(3), 731–751.

Hausch, D. B., & Li, L. (1993). A common value auction model with endogenous entry and information acquisition. *Economic Theory*, *3*, 315–334.

International Renewable Energy Agency. (2019). *Renewable energy auctions: Status and trends beyond price (preliminary findings).* Retrieved from `www.irena.org/-/media/Files/IRENA/Agency/Publication/2019/Jun/IRENA_Auctions_beyond_price_2019_findings.pdf`

Kim, K., & Koh, Y. (2022). Auctions with flexible information acquisition. *Games and Economic Behavior*, *133*, 256–281.

Kleinberg, R., Waggoner, B., & Weyl, E. G. (2018). Descending price optimally coordinates search. *Available at SSRN: https://ssrn.com/abstract=2753858*.

Klemperer, P. (2002). What really matters in auction design. *Journal of Economic Perspectives*, *16*(1), 169–189.

Levin, D., & Smith, J. L. (1994). Equilibrium in auctions with entry. *American Economic Review*, 585–599.

Lu, J., Ye, L., & Feng, X. (2021). Orchestrating information acquisition. *American Economic Journal: Microeconomics*, *13*(4), 420–465.

McAfee, R. P., & McMillan, J. (1987). Auctions with entry. *Economics Letters*, *23*(4), 343–347.

Miettinen, P. (2013). Information acquisition during a Dutch auction. *Journal of Economic Theory*, *148*(3), 1213–1225.

Pavan, A., Segal, I., & Toikka, J. (2014). Dynamic mechanism design: A Myersonian approach. *Econometrica*, *82*(2), 601–653.

Persico, N. (2000). Information acquisition in auctions. *Econometrica*, *68*(1), 135–148.

Quint, D., & Hendricks, K. (2018). A theory of indicative bidding. *American Economic Journal: Microeconomics*, *10*(2), 118–51.

Rezende, L. (2018). Mid-auction information acquisition. *Economic Theory*, *65*(3), 751–780.

Roberts, J. W., & Sweeting, A. (2013). When should sellers use auctions? *American Economic Review*, *103*(5), 1830–1861.

Shi, X. (2012). Optimal auctions with information acquisition. *Games and Economic Behavior*, *74*(2), 666–686.

Stegeman, M. (1996). Participation costs and efficient auctions. *Journal of Economic Theory*, *71*(1), 228–259.

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica*, *47*(3), 641–654.

Ye, L. (2007). Indicative bidding and a theory of two-stage auctions. *Games and Economic Behavior*, *58*(1), 181–207.