CESifo WORKING PAPERS

7394
2018

December 2018

Selling Complementary Goods: Information and Products

Suehyun Kwon

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl
www.cesifo-group.org/wp

An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: www.CESifo-group.org/wp

CESifo *Working Paper No. 7394*
Category 11: Industrial Organisation

Selling Complementary Goods: Information and Products

Abstract

This paper studies optimal mechanisms for selling complementary goods sequentially. The seller starts with private information, has limited commitment and offers in the first period a menu of information structures on the value of the second-period product. Fully revealing the seller type in the first period makes the second period a standard adverse selection problem, and fully revealing the buyer type in the first period makes the second period an information design problem. Among properties of equilibria, all types of seller must pool in every equilibrium if certain first-order stochastic dominance and independence conditions are satisfied.

Keywords: information design, dynamic informed-principal problem, interdependent values, limited commitment, Myerson-Satterthwaite.

Suehyun Kwon
Department of Economics
University College London
30 Gordon Street
United Kingdom – London, WC1H 0AX
suehyun.kwon@gmail.com

November 27, 2018
First draft: August 15
This paper was previously circulated as “Pricing Advices.” I thank Joel Sobel and seminar participants at Yale for their comments.

1 Introduction

Many financial institutions offer advisory services. Sometimes, clients only pay for advice, but typically these institutions also offer financial products that they recommend at the end of an advisory service. The client's willingness to pay for a financial product or whether he is going to purchase any financial product in the first place depends on the advice he gets. This paper studies how to price when a seller charges for complementary goods sequentially, and particularly, what happens if the first product is information about the value of the second product.

I model this problem as designing a selling mechanism where the seller offers a menu of information structures in the first period and a menu of (price,quantity)- or (price,quality)-schedule in the second period. Neither the seller nor the buyer has to commit to any period. The buyer can choose whether to participate in the first period then chooses an information structure; after observing the signal privately, the buyer decides whether to purchase in the second period. The buyer can also choose only to participate in the second period without purchasing an information structure.

Apart from prices the seller charges, ex-post payoffs of the seller and the buyer depend on the payoff-relevant state (state of the world) and whether trade happens in the second period. I allow both the seller and the buyer to have private information, and this allows for interdependent values. The payoff-relevant state could also be the pair of (seller type, buyer type); the seller's private information is his type, and the buyer's private information is his type. This includes private values as in Myerson-Satterthwaite no-trade theorem (1983). Alternatively, both the seller and the buyer can have informative signals about the payoff-relevant state which is the buyer's valuation of the product. This maps into the informed-principal problem, and the typical assumption in any monopoly setting is that the seller's payoff doesn't depend on the payoff-relevant state.

Allowing for interdependent values and informed-principal problems leads to many real-world applications. As mentioned already, financial advisory services or private banking industry offer advices as a core part of their business. Many skincare brands offer some type of analysis or advice, and sometimes these are for a fixed fee, but sometimes the advice is offered for free. There are also educational advisory services in many different countries including the US which help students prepare for college entrance exams or college applications. In this case, one could think of the seller offering consultation then preparation for college. Whenever the buyer or the client has to first get an estimate for the service, the seller offers the information for a fee then the product or the service if the buyer decides to purchase. The information disclosure literature has focused on the case where this information is offered for free, and there is a significant difference if the seller can charge for information and the buyer can choose which period to participate in the mechanism if he is going to participate at all.

When the seller offers a menu of information structures, the seller first charges a price for each information structure, then a signal is realized. The probability of each signal is known when the buyer decides whether to purchase the information structure. I assume the seller doesn't observe the signal realization. An information structure can depend both on the seller's type and the payoff-relevant state. If the seller has no private information and/or the information structure is a mapping from the set of payoff-relevant states to the set of signals, this allows for experiments as in the Bayesian persuasion literature.

By offering an information structure, the buyer in most cases won't have complete information even after purchasing the information structure. However, by allowing information structures to be conditional both on the seller's type and the payoff- relevant state, the seller can sell a signal informative about his own private information and also informative about the payoff-relevant state that he doesn't know himself. Since the seller can have private information and sells both an information structure and a product, there are a few branches of literature this paper can be related to in principle; however, there are very few papers on these. The related literature includes (i) dynamic informed-principal problem, (ii) informed-principal selling information,

(iii) selling both the information and the product, or more generally, (iv) selling complementary products. It is also related to (v) mechanism with limited commitment. Complementarity between the information structure and the product is related to common agency, but in my model, the single seller sells both. In terms of a single seller selling complementary products, this generally involves multi-dimensional screening, and there are very few papers both in the simultaneous pricing case and the sequential pricing case.[1] I will first describe the results then list the few existing related papers.

Key features of my model are (i) two-sided private information allowing for interdependent values, (ii) limited commitment, and (iii) the first-period product being information about the value of the second-period product. These lead to methodological challenges since some of the usual tools cannot be applied. Qualitatively, (i)-(iii) lead to the following observation; if the seller fully reveals his type in the first period, then the buyer is the only player with private information in the second period. Therefore, the second period turns into a standard adverse selection problem. I assume the Spence-Mirrlees condition which implies that there will be a unique optimal mechanism in the second period in this case. On the other hand, if the buyer fully reveals his type in the first period, the seller is the only player with private information in the second period. Now, this can be compared to the information design literature in which the designer provides a partition of his information to a sequence of short-run agents. The second period after the buyer fully reveals his type in the first period maps into a short-run agent in an information-design problem without loss of generality as the seller can just post a price. All of (i)-(iii) are necessary for these properties.

As for pooling strategies, I provide a sufficient condition for every equilibrium to have all types of seller to pool in both periods. This is a condition on the primitives of the model, i.e., the common prior π on (state of the world, seller type, buyer type) and the set of information structures allowed for the seller. Loosely speaking, when both the seller and the buyer have limited commitment, one has to characterize equilibria of the extensive-form game,

[1] Daskalakis et al (2017) has additively separable utility.

and these conditions guarantee that in any equilibrium, the seller can rank all menus of information structures and all menus of contracts that can be offered with a strictly positive probability for some seller type. Since for each seller type, the actions are ordered identically, no seller type will choose any strictly dominated action. It might sound obvious, but this condition is satisfied with Bayesian persuasion, i.e., the information structure independent from each seller type, and first-order stochastic dominance.

I certainly do not characterize every equilibrium allowing for all correlations in the common prior π. One can also characterize sufficient conditions for offering a single piece of information vs. creating endogenous types to be optimal, in which case the buyer's consumption utility in the second period matters. When the information structure only depends on the state of the world and not on the seller type, the buyer's posterior is a martingale after purchasing an information structure. However, the buyer's effective type in the second period depends on his consumption utility in each state of the world, and if this is linear, then the buyer's effective type in the second period is also a martingale of his expected type before purchasing any information structure in the first period. This no longer holds if the buyer's utility is non-linear.

As for methodological challenges, one can ask what happens to the revelation principle and the one-stage deviation principle in this class of models. This leads to the discussion on endogeneity of learning process which will be further discussed in the conclusion. Informed principal with limited commitment makes it difficult to extend Bester-Strausz (2001) in a straightforward manner by just modifying the buyer's reporting stragety given $\mathcal{M} = \Omega_b$. (or $\Omega_b \times S$ in the second period) When the buyer deviates in the first period and chooses an information structure that his true type chooses with zero probability on the equilibrium path, then the seller assigns zero probability to the buyer's true type in the second period, and given a menu of contracts, the buyer has no bundle intended for his true type and the signal that was realized at the end of the first period. Compared to the one-stage deviation principle in the usual sense in mechanism design literature, there is no "on-the-equilibrium-path" action that the seller is expecting the buyer to take, but this has to be discussed together with the revelation principle. Given that in both periods, the seller offers a menu of information structures or a menu of contracts, if one were to apply Bester-Strausz (2001) literally, one might think that the seller can ask the buyer to report his private information in the beginning of both periods, but the buyer might not report truthfully in both periods with every buyer type realization. In my model, given that there are only two periods, one can do backward induction, and invoking the revelation principle is not a necessity. However, this points at a bigger problem on the revelation principle once we move away from the most standard mechanisms, and the issue is worth pointing out.

As mentioned already, all three points (i)-(iii) matter for the revelation principle. Bester-Strausz (2001) has a single agent with a finite time horizon, and the utility from allocations only matters within each period; the mechanism designer has limited commitment but doesn't have any private information himself. There are examples in the literature, (for example, see Pavan-Toikka-Segal (2013) or Kwon (2018)) where the one-stage deviation principle in the usual sense is without loss of generality if the environment is first-order Markov, but these papers mostly assume full commitment of the mechanism designer. My model has two periods, but the environment is not a first-order Markov process; the first-period product is information on the value of the second-period product, and if one were to consider ω_b as the buyer type in the beginning of the first period, and (ω_b, s) as the buyer type in the beginning of the second period, then the mapping is between Ω_b to $\Omega_b \times S$. The idea behind the one-stage deviation principle in existing papers is not that the type space never changes over time, but that the current type is a sufficient summary statistics of the agent's private history up to that point. One could take the buyer's posterior on the joint distribution of $\Omega \times \Omega_s$, but this approach can be justified only if there is no other relevant information, and in order to detect any deviation, the buyer needs to keep track of the support of each action of the seller and the signal realization. If one were to construct a proof by (i) finding a reporting strategy of the buyer when the seller has full commitment and (ii) finding an equivalent strategy when the seller has limited commitment, then we cannot prove the one-stage deviation principle for the full-commitment solution in the first place.

The rest of the paper is organized as follows. I will discuss remaining related papers then describe the model in Section 2. Results are in section 3. Section 4 concludes.

1.1 Related Literature

The most relevant literature for this paper is information design and information disclosure. Three key assumptions of the model, two-sided private information, limited commitment, and the first-period product being information on the second-period product, distinguish results in this paper from the information design literature and the information disclosure literature. However, my results should be aligned in the literature for the mechanism designer providing "information."

Both the information design literature and the information disclosure literature have more recent papers than what can be attributed to these two topics from decades ago. Kremer, Mansour and Perry (2014) and Che, Hörner (2018) study a sequence of myopic agents with an information designer. Even though I have only two periods in this model, the agent or the buyer doesn't have any commitment power, and if the buyer reveals his type in the first period, then the seller is the only one with private information in the second period, and the buyer's decision problem after the seller offers a menu of contracts is identical to the myopic agents' in these papers. However, the buyer in my model need not disclose his type fully in the first period.

With information disclosure, one could trace it back to Crawford-Sobel (1982) or Milgrom (1981), depending on whether the information is soft or hard. Li-Shi (2017) considers a seller disclosing information about the product before selling it, and together with Bergemann-Pesendorfer (2007) and Eso-Szentez (2007), these are more closely related to my paper. However, limited commitment makes a difference in my model, and the most immediate difference is that the seller cannot fully extract the ex-ante expected rent from the buyer by charging it as the first-period price of the information structure in my model. This argument only works with full commitment.

Bayesian persuasion literature starts with Kamenica-Gentzkow (2011), and the difference between Bayesian persuasion and information design in general is whether the designer knows the information he provides. Some papers in the information disclosure literature assume either of the two possibilities, but the mechanism designer in the Bayesian persuasion literature offers an experiment on the state of the world, which he doesn't know himself either.

As for the key assumptions of the model, see Mylovanov-Troeger (2014) for the informed-principal problem. Given that my model has only two periods, even though limited commitment does play a key role in most results, one should think of it more as PBE of extensive-form game. Monopoly pricing with complementarity also has not been studied extensively, and I am not aware of any publication studying complementarity of information and the product itself. For common agency, refer to works of Pavan.

Most of this paper focuses on revenue maximization where the seller has zero reservation utility; for non-zero reservation utility, there are Garratt-Pycia (2016), Roesler-Szentes (2017) and Condorelli-Szentes (2018).

Some of the learning dynamics in my paper is related to common-value auctions and winner's curse, but most existing papers focus on interdependent values or correlated information across different players and don't characterize what happens if the designer's information or payoff is correlated with the agent's. For auctions with interdependent values, see Kojima (2017).

2 Model

There are one seller and one buyer. The seller has a product, and both the seller's type, ω_s, and the buyer's type, ω_b, are private information. The payoff-relevant state is ω. The seller and the buyer first learn their types then the seller offers a mechanism; the mechanism charges for an information structure in the first period and for a product in the second period. In the second period, the seller can offer a price-quantity schedule or a price-quality schedule, (q, p).

I focus on limited commitment with both the seller and the buyer. Both the seller and the buyer have quasilinear utilities, and there's no discounting between two periods. Let p_1, p_2 be prices in the first and the second period, respectively. If the buyer doesn't participate in period t, denote $p_t = 0$. The seller's payoff is $u_s(q|\omega) + p_1$ if he doesn't sell the product in the second period, and $p_1 + p_2$ if he sells the product. The buyer's payoff is $u_b(q|\omega) - p_1 - p_2$ if he buys the product and $-p_1$ if he doesn't buy the product.

This setup allows for different types of private information. If the payoff-relevant state is the pair of (seller type, buyer type), we have $\omega = (\omega_s, \omega_b)$. A special case is $u_s(q|\omega) = \omega_s$, $u_b(q|\omega) = \omega_b$; the seller and the buyer know their own valuations of the product, and there is one unit to be sold. This maps into two-sided private information and Myerson-Satterthwaite (1983). Alternatively, we can have $u_s(q|\omega) = 0$ as in the usual seller-buyer setting, and both the seller and the buyer have informative signals about the buyer's valuation of the good which is the payoff-relevant state ω. This setup is related to informed-principal problems and the information disclosure literature as in Bergemann-Pesendorfer (2007), Eso-Szentes (2007) and Li-Shi (2017). Or both $u_s(q|\omega)$ and $u_b(q|\omega)$ depend on the payoff-relevant state, and this allows for interdependent values. It is without loss of generality to define the utility function as a function of signals, but once the buyer acquires information, we need to make more assumptions to define how it changes the expected utility of the buyer from purchasing the good in the second period.

$\omega_s \in \Omega_s$, $\omega_b \in \Omega_b$, $\omega \in \Omega$ and $\Omega_s, \Omega_b, \Omega$ are non-empty metric spaces. The common prior over $\Omega \times \Omega_s \times \Omega_b$ at the beginning of the game is π. Since Ω_s, Ω_b are not necessarily real numbers, I don't make any assumptions on u_s, u_b for now. The current formulation $u_s(q|\omega), u_b(q|\omega)$ only depends on ω, and essentially, the seller or the buyer needs to take an expectation of $\mathbb{E}[u_s(q|\omega)|\omega_s], \mathbb{E}[u_b(q|\omega)|\omega_b]$ in the beginning of the first period. If after the first period, after the seller offers the menu of information structures and after the buyer chooses whether to participate and which information structure to purchase, both the seller and the buyer can learn about each other's private information; denote the probability distribution over Ω_s revealed by the menu of information structures as π_s and the probability distribution over Ω_b revealed by the buyer's choice as π_b. For π_s, it only signals the seller's private information, and π_s doesn't depend on ω_b. For π_b, when the buyer decides to participate or chooses an information structure, he already knows both ω_b and π_s; the buyer should update from the joint prior π, and then his choice signals ω_b to the seller. The Bayesian updating is further discussed in the next section, but this allows for ω_s, ω_b to be information on the payoff-relevant state, a part of the payoff-relevant state, a combination of both and so forth. Depending on the further specification, there will be restrictions on the common prior π, for example, when ω_s, ω_b are part of ω.

The type of information structure the seller can offer in the first period is $\gamma : \Omega_s \times \Omega \to \Delta(S)$ where S is a metric space and $\Delta(S)$ is the set of distributions on S. Since the seller learns his type before offering a mechanism, the seller can charge different prices conditional on his type. When Ω_b is not a singleton, the buyer has private information, and the seller offers a menu of information structures that the buyer can choose from; by the revelation principle, each buyer type has an information structure intended for his type. If the buyer chooses an information structure, the buyer observes the signal privately. Afterwards, the seller offers a price-quantity schedule or a price-quality schedule in the second period.

The information structure the seller can offer in the first period can depend both on the seller's private information and the payoff-relevant state. A special case of this information structure is Bayesian persuasion where the seller has no private information and offers an experiment without knowing the signal realization nor the payoff-relevant state. More precisely, if the information structure is a mapping $\gamma : \Omega \to \Delta(S)$, then the seller can provide an additional signal about the payoff-relevant state that is independent of his own signal. This implies that the seller can offer information that he doesn't know himself. I assume the signal realization of the information structure is independent of the buyer's private information. I also consider another special case of information structure where $\gamma : \Omega_s \to \Delta(S)$. In this case, the seller can only offer information about his own private information. This case is irrelevant when the seller has no private information, and the relevant cases are when only the seller has private information or when both the seller and the buyer have private information.

Assumption 1 is a standard assumption for static adverse selection problems, and this could be generalized to the usual Spence-Mirrlees condition. This assumption maps Ω to $\mathbb{R}_+$, and it might seem like a one-dimensional type space, but any prior on the state of the world maps into expected utility and has the same property. One should also note that the multiplicative separability is between ω and q, and there is no hidden assumption on Ω, Ω_s and Ω_b.

Assumption 1. *The buyer's payoff in the second period from purchasing* (q, p) *is* $u_b(p, q|\omega) = U(\omega)V(q) - p$ *where* $V(0) = 0,\ V'(\cdot) > 0 > V''(\cdot),\ q \in \mathbb{R}_+$ *is the quantity or quality, and* $p \in \mathbb{R}+$ *is the price. Further assume* $U : \Omega \to \mathbb{R}_+$ *and define* $\mathbb{R}_+ = \{x|x \in \mathbb{R}, x \geq 0\}$*. The seller's payoff is* $p - cq$ *which requires a constant marginal cost.*

Throughout the rest of the paper, a menu of information structures is denoted by Γ, each information structure by γ, and each signal by s. The menu of contracts in the second period is Γ_2. μ denotes the seller's strategy at each point, and σ denotes the buyer's strategy. With limited commitment, the solution concept is Perfect Bayesian equilibrium.

3 Results

Preliminary results in section 3.1 explain Bayesian updating in this model a bit more carefully. Earlier versions of this paper had results when at most one of the seller or the buyer has private information, but they are not part of the current draft. After section 3.1, the following two subsections characterize optimal mechanisms when both parties have private information. In section 3.2, I assume that the seller doesn't value the good himself in the second period as in the usual monopoly setting, i.e., $u_s(q|\omega) \equiv 0$. Section 3.3 discusses what happens when $u_s(q|\omega) \neq 0$ which maps into interdependent values and also private values as in Myerson-Satterthwaite (1983). $u_s(q|\omega) \neq 0$ makes the seller's commitment power more relevant, but detailed characterization with the seller's reservation utility is deferred to a different paper.

3.1 Preliminary Results

In section 2, all three spaces, $\Omega, \Omega_s, \Omega_b$ are assumed to be any metric spaces, but for the rest of the paper, one can read any integration as Lebesgue integration; these spaces are equipped with Lebesgue measure.

I will next derive the Bayesian updating throughout the first period. At the beginning of first period, the common prior on $\Omega \times \Omega_s \times \Omega_b$ is π, and both the seller and the buyer have private information ω'_s, ω'_b. With a slight abuse of notation, denote the marginals from π as $\pi(\omega, \omega_b|\omega_s)$ and $\pi(\omega, \omega_s|\omega_b)$.

$$\begin{aligned}
\pi(\bar{\omega}, \bar{\omega}_b|\omega'_s) &= \frac{\pi(\bar{\omega}, \bar{\omega}_b, \omega'_s)}{\int \pi(\omega, \omega_b, \omega'_s) d\omega d\omega_b}, \\
\pi(\bar{\omega}, \bar{\omega}_s|\omega'_b) &= \frac{\pi(\bar{\omega}, \omega'_b, \bar{\omega}_s)}{\int \pi(\omega, \omega'_b, \omega_s) d\omega d\omega_s}.
\end{aligned}$$

Then the seller offers a menu of information structures, Γ, and with each type of seller offering a particular menu of information structures with probability $\mu(\Gamma|\omega_s)$, the buyer can update his posterior from $\pi(\omega, \omega_s|\omega'_b)$ to

$$\pi(\bar{\omega}, \bar{\omega}_s|\omega'_b, \Gamma) = \frac{\pi(\bar{\omega}, \bar{\omega}_s|\omega'_b)\mu(\Gamma|\bar{\omega}_s)}{\int \pi(\bar{\omega}, \omega_s|\omega'_b)\mu(\Gamma|\omega_s) d\omega_s}.$$

Given this, now the buyer decides whether to participate, and if he does, which information structure to purchase. Afterwards, the seller updates his belief once again, but this step of Bayesian updating requires the seller having a belief over $\pi(\omega, \omega_s|\omega'_b, \Gamma)$ then updating once more and requires a bit more care than other steps.

When the seller offers Γ, the seller can update his belief about the buyer's marginal $\pi(\omega, \omega_s|\omega'_b, \Gamma)$, but the seller still only knows his own private information ω_s and his strategy $\mu(\Gamma|\omega_s)$. Given the common prior π, the set of menus of information structures that some type of seller offers with a strictly positive probability, the set of $\pi(\omega, \omega_s|\omega_b', \Gamma)$ is already given, and the seller's belief is given by his marginal $\pi(\omega, \omega_b|\omega_s')$.

The key here is that the buyer can update his belief from the common prior to the marginal given his own private information then once again given the menu of information structures, and it is reflected in his marginal $\pi(\omega, \omega_s|\omega_b', \Gamma)$, but from the seller's point of view, he only knows the menu of information structures Γ he offered. Furthermore, once the seller considers the set of $\pi(\omega, \omega_s|\omega_b', \Gamma)$ given Γ, the seller just needs to know ω_b', and given that the buyer hasn't taken any action up until that point, the seller cannot update any further from the marginal given his own private information ω_s.

However, when the buyer chooses an information structure or decides not to participate, the seller can update his belief once more. At this stage, let $\sigma(\gamma|\omega_b, \Gamma)$ be the buyer's strategy given his private information ω_b and the menu of information structures Γ. γ encodes no participation as an option in addition to information structures in Γ.

$$\pi(\bar{\omega}, \bar{\omega}_b|\omega_s', \Gamma, \gamma) = \frac{\pi(\bar{\omega}, \bar{\omega}_b|\omega_s')\sigma(\gamma|\bar{\omega}_b, \Gamma)}{\int \pi(\bar{\omega}, \omega_b|\omega_s')\sigma(\gamma|\omega_b, \Gamma)d\omega_b}.$$

The seller knows which Γ he offered, but given that the buyer's strategy depends on the menu of information structures the seller offered, the marginal has to keep track of Γ as well.

After choosing γ and observing a particular signal s, the buyer can update his posterior belief from

$$\pi(\bar{\omega}, \bar{\omega}_s|\omega_b', \Gamma) = \frac{\pi(\bar{\omega}, \bar{\omega}_s|\omega_b')\mu(\Gamma|\bar{\omega}_s)}{\int \pi(\bar{\omega}, \omega_s|\omega_b')\mu(\Gamma|\omega_s)d\omega_s}$$

to

$$\pi(\bar{\omega}, \bar{\omega}_s|\omega_b', \Gamma, \gamma, s) = \frac{\pi(\bar{\omega}, \bar{\omega}_s|\omega_b', \Gamma)\mu(s|\gamma, \bar{\omega}, \bar{\omega}_s)}{\int \pi(\omega, \omega_s|\omega_b', \Gamma)\mu(s|\gamma, \omega, \omega_s)d\omega d\omega_s}.$$

After Γ_2, we get

$$\pi(\bar{\omega}, \bar{\omega}_s|\omega_b', \Gamma, \gamma, s, \Gamma_2) = \frac{\pi(\bar{\omega}, \bar{\omega}_s|\omega_b', \Gamma, \gamma, s)\mu(\Gamma_2|\bar{\omega}_s, \Gamma, \gamma)}{\int \pi(\bar{\omega}, \omega_s|\omega_b', \Gamma, \gamma, s)\mu(\Gamma_2|\omega_s, \Gamma, \gamma) d\omega_s}.$$

We can also specify the seller's belief after the buyer chooses γ. Since the buyer observes s privately, the seller only knows the distribution of s which is independent of ω_b.

So far, Bayesian updating itself only requires the common prior π and strategies $\mu(\Gamma|\omega_s), \sigma(\gamma|\omega_b, \Gamma)$. In any equilibrium, these have to be best responses.

3.2 Revenue Maximization: $u_s(q|\omega) = 0$

When $u_s(q|\omega) = 0$, the seller doesn't value keeping the good to himself and wants to maximize the revenue from trade. The main difference between this section and the next section $u_s(q|\omega) \neq 0$ is that when the seller's valuation from keeping the good himself depends on the payoff-relevant state, the seller might not always want to trade, and the buyer's private information is informative about whether the seller prefers to trade; otherwise, the seller always prefers to sell and maximize the revenue. When $u_s(q|\omega) = 0$, the seller wants to maximize revenue which implies that it's the buyer's willingness to pay that depends on private information of the seller and the buyer; this section is closely related to the informed-principal problem.

Results in this section differ from other mechanisms with two periods in a sense that the first period allocation, i.e., the information structure, changes the type distribution in the second period. In order to compare full-commitment solutions to limited-commitment solutions, it will help to focus on full-commitment solutions where the second-period mechanism is optimal. However, this may not always be without loss of generality, and I do not characterize full-commitment solutions in this paper. I will discuss the difference between other types of allocation to information as I characterize relevant theorems. Furthermore, it is not obvious a priori that the revelation principle extends to limited-commitment solutions with informed principal.

The first lemma characterizes the standard optimal static mechanism in the second period. The only difference is that the payoff-relevant state is multiplicatively separable from the quantity/quality, but one can order $\mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma, s]$.

Lemma 1. *In the second period, the optimal menu of contracts is given by the standard results with Spence-Mirrlees condition: for given Γ and γ, order $\Omega_b \times S$ by $\mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma, s]$. Let $\Pr(s|\gamma, \omega, \omega_s)$ be the probability of signal s in the information structure γ, and $F(\omega_b, s|\omega_s, \Gamma, \gamma)$ is the corresponding CDF, when (ω_b, s) are ordered by $\mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma, s]$. The pdf is $f(\cdot)$. When there is no further inference based on the second-period menu of contracts, the optimal menu is no trade below the cutoff type (ω_b', s'), and otherwise,*

$$(\mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma, s] - \frac{1-F(\omega_b, s)}{f(\omega_b, s)})V'(q(\omega_b, s)) = c,$$
$$T(\omega_b, s) = \mathbb{E}[U(\omega)|\omega_b, s]V(q(\omega_b, s)) - \int_{(\omega_b', s')} V(q(\omega_b, s))d\tau,$$

where $\tau(\omega_b, s)$ is the indicator function whether the seller sells to that type. The seller's expected revenue is

$$\int_{(\omega_b', s')} ((\mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma, s] - \frac{1-F(\omega_b, s)}{f(\omega_b, s)})V(q(\omega_b, s)) - cq(\omega_b, s))dF.$$

Proof. In the second period, we already know that the buyer's utility satisfies Spence-Mirrlees condition, and the usual adverse selection results hold. The cutoff type after observing the signal, (ω_b', s') gets zero rent, and the optimal contract is pinned down by the local IC. Given the one-dimensional ordering on $\Omega_b \times S$, we can take $\theta = (\omega_b, s)$ as the type of the buyer. Then the usual results apply and we have

$$\theta V'(q(\theta)) = c + \frac{1-F(\theta)}{f(\theta)}V'(q(\theta)) \;\Leftrightarrow\; (\theta - \frac{1-F(\theta)}{f(\theta)})V'(q(\theta)) = c$$
$$T(\theta) = \theta V(q(\theta)) - \int_{\underline{\theta}}^{\theta} V(q(\tau))d\tau$$

$$\Rightarrow \mathbb{E}[U(\omega)|\omega_b,\Gamma,\gamma,s]V'(q(\omega_b,s)) = c + \frac{1-F(\omega_b,s)}{f(\omega_b,s)}V'(q(\omega_b,s))$$
$$\Leftrightarrow (\mathbb{E}[U(\omega)|\omega_b,\Gamma,\gamma,s] - \frac{1-F(\omega_b,s)}{f(\omega_b,s)})V'(q(\omega_b,s)) = c$$
$$p(\omega_b,s) = \mathbb{E}[U(\omega)|\omega_b,\Gamma,\gamma,s]V(q(\omega_b,s)) - \int_{(\omega_b',s')} V(q(\omega_b,s))d\tau$$

where $\tau(\omega_b,s)$ is the indicator function whether the seller sells to that type.

The seller wants to sell to (ω_b,s) if and only if the virtual surplus is weakly positive for those types: the seller's expected payoff from the second period is

$$\int_{(\omega_b',s')} (p(\omega_b,s) - cq(\omega_b,s))dF$$
$$= \int_{(\omega_b',s')} (\mathbb{E}[U(\omega)|\omega_b,\Gamma,\gamma,s]V(q(\omega_b,s)) - \int_{(\omega_b',s')} V(q(\omega_b,s))d\tau - cq(\omega_b,s))dF$$
$$= \int_{(\omega_b',s')} ((\mathbb{E}[U(\omega)|\omega_b,\Gamma,\gamma,s] - \frac{1-F(\omega_b,s)}{f(\omega_b,s)})V(q(\omega_b,s)) - cq(\omega_b,s))dF$$

and s is ordered by $\mathbb{E}[U(\omega)|\omega_b,\Gamma,\gamma,s]$. □

Lemma 1 is a standard result, but the specification of the model allows for a natural one-dimensional ordering of $\Omega_b \times S$. In principle, I can allow S to be any metric space, and different types of buyer may have different ordering and the cutoff on S, conditional on their type. One can characterize sufficient conditions when they coincide; one such case is when ω_b is a prior and each s can be ordered by the first-order stochastic dominance. But it is unnecessary for the rest of the paper. One should also note that $F(\omega_b,s|\omega_s,\Gamma,\gamma)$ is the seller's belief on the joint distribution of (ω_b,s); the buyer knows his own ω_b and privately observes s.

Now before characterizing an optimal mechanism, one needs to discuss what it means when an information structure is a mapping from $\Omega_s \times \Omega$ to $\Delta(S)$. When there is an information structure on Ω and the seller doesn't know the true payoff-relevant state, then the posterior beliefs after observing the signal realization should be a martingale. When the seller can offer an informative signal about his own private information, then one can think about what types of informative signals the seller can provide. In particular, what should be the difference between an informative signal based on the private information of the seller and an inference based on the menu of information structure as in the usual informed-principal problem? In the information design literature, information is a partition of the designer's information most of the times. Also in any application, an equilibrium as in Myerson requires that the buyer should be able to make the inference. To be consistent with three branches of literature, I assume (i) the information structure partitions the seller's information, then (ii) adds a martingale, (iii) the buyer can still make inferences given the menu of information structures. One can allow for more general information structures than in (i), but I will restrict the set of available information structures in this paper; this makes it consistent with information design literature. Most importantly, (iii) means that if there is only one type of seller offering a particular menu of information structures, then the seller reveals his private information.

One also needs to discuss the difference between full commitment and limited commitment on the seller side. With limited commitment, the mechanism designer will always offer an optimal one in the second period, but with full commitment, it may not be optimal to offer an ex-post optimal menu of contracts in the second period. Also, when there are two periods, one needs to think about the buyer's incentives for double deviation, and offering a menu of information structures in the first period has nontrivial effects on double deviation in addition to revelation principle. For example, lemma 1 shows that the marginal type (ω_b', s') depends on each buyer type who pools together in the first period. If the buyer chooses an information structure his type chooses with 0 probability on the equilibrium path, after the buyer observes s, we need to compare (ω_b, s) to (ω_b', s') for all ω_b' who purchases that information structure with a strictly positive probability on the equilibrium path. From lemma 1, it's not just the effective type $\mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma, s]$ that is affected by the first-period deviation. This expectation can actually be taken care of by finding the equivalent types $(\tilde{\omega}_b, \tilde{s})$. However, now the local IC in the usual adverse selection model with Spence-Mirrlees condition shows that the marginal type for each information structure depends on the buyer types who purchase on the equilibrium path, and the seller needs to consider the rent each buyer type can get by deviating and taking the bundle for the effectively equivalent type. This makes it easier to solve for the limited commitment case which is in the next theorem.

Theorem 1 (Limited Commitment)**.** *There exists no profitable and undetectable deviation for either the seller or the buyer if the equilibrium strategies maximize the following two objective functions simultaneously:*

$$\max_{\Pi(\Gamma)} \Bigg(\int_{\gamma\in\Gamma} p_1(\gamma)\sigma(\gamma|\omega_b,\Gamma)d\pi(\omega_b|\omega_s) \\ + \max_{\Pi(\Gamma_2)} \int_{\gamma_2\in\Gamma_2} p_2(\gamma_2)\gamma(s|\omega,\omega_s)\frac{\pi(\omega,\omega_b|\omega_s')\sigma(\gamma|\omega_b,\Gamma)\sigma_2(\gamma_2|\omega_b,\Gamma,\gamma,s,\Gamma_2)}{\int \pi(\omega,\bar{\omega}_b|\omega_s')\sigma(\gamma|\bar{\omega}_b,\Gamma)d\bar{\omega}_b} d\omega d\omega_b \Bigg)$$

$$\max_{\gamma\in\Gamma} \Bigg(-p_1(\gamma) + \int \max_{\gamma_2\in\Gamma_2} \Big(\int U(\bar{\omega}) \Big(\frac{\pi(\bar{\omega},\bar{\omega}_s|\omega_b')\mu(\Gamma|\bar{\omega}_s)\gamma(s|\bar{\omega},\bar{\omega}_s)\mu(\Gamma_2|\bar{\omega}_s,\Gamma,\gamma)}{\int \pi(\omega,\omega_s|\omega_b')\mu(\Gamma|\omega_s)\gamma(s|\omega,\omega_s)\mu(\Gamma_2|\omega_s,\Gamma,\gamma)d\omega d\omega_s} \Big) d\bar{\omega}d\bar{\omega}_s \\ \times V(q(\gamma_2)) - p_2(\gamma_2) \Big) \mu_2(\Gamma_2|\tilde{\omega}_s,\Gamma,\gamma) \\ \int \frac{\pi(\tilde{\omega},\tilde{\omega}_s|\omega_b,\Gamma,\gamma)\gamma(s|\tilde{\omega},\tilde{\omega}_s)}{\int \pi(\omega,\omega_s|\omega_b,\Gamma,\gamma)\gamma(s|\omega,\omega_s)d\omega d\omega_s} d\tilde{\omega} \Big) d\tilde{\omega}_s \Bigg).$$

Any omitted proofs including the proof of theorem 1 are in the appendix. I will discuss the IC constraints then move on to the properties of PBE. Also the only detectable deviations in this model are either the seller offering an off-the-equilibrium-path menu or the buyer's participation decisions; one does need to keep track of the support of each action on the equilibrium path, and this is discussed in more detail in the next few paragraphs.

Lemma 1 is the unique optimal menu of contracts in the second period given $F(\omega_b, s|\omega_s, \Gamma, \gamma)$ if neither the buyer nor the seller deviated in the first period. However, when the seller offers a second-period menu of contracts, Γ_2 signals ω_s. As for the buyer's decision in the second period after the menu of contracts is offered, he is just maximizing his expected utility. One can think of the rest of the second period after a menu of contracts is offered as (i) the buyer updates his posterior to $\pi(\omega, \omega_s|\omega_b, \Gamma, \gamma, s, \Gamma_2)$, (ii) the buyer chooses a bundle if he participates. Given that the buyer's effective type $\mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma, s, \Gamma_2]$ doesn't depend on ω_s directly, this only affects the inference of s and the first-period information structure that the buyer already purchased. This is where the correlation across $\Omega, \Omega_s, \Omega_b$ matters; compared to static infomed-principal problems, now Γ_2 can change the inference of the previous period. Along the equilibrium path, Γ gives a set of seller types who offers this menu together with conditional probabilities. γ is just a part of Γ that is already offered. The probability of each s depends on γ and when it is a mapping from both Ω and Ω_s, depends on ω_s as well. On the equilibrium path, Γ_2 after Γ, γ are observed by both the seller and the buyer again gives a set of seller types who offer this menu together with conditional probabilities.

If a seller type ω_s reveals his type by the time he offers Γ_2, then the buyer has his posterior on ω given $\omega_s, \omega_b, \Gamma, \gamma, s, \Gamma_2$. The seller's posterior on this buyer's posterior, or effectively $\mathbb{E}[U(\omega)|\omega_s, \omega_b, \Gamma, \gamma, s, \Gamma_2]$, characterizes the ex-post optimal mechanism in the second period. As long as any two ω_s with the same marginals on $\Omega \times \Omega_b$ don't pool in the first period, the seller can reveal his type by offering the ex-post optimal mechanism in the second period; but this only shows existence of such equilibria, and if the seller deviates in a PBE, it depends on the buyer's off-the-equilibrium-path belief what he thinks is the seller's type.

If the seller doesn't reveal his type, incentives due to informed-principal problem can be summarized as follows. Once Γ_2 is offered, the buyer has his posterior π given all available information up to that point, and the buyer makes his decision. One can consider when Γ_2 is ex-post optimal given the updated posterior of the buyer after observing Γ_2. But this need not be the case. On the equilibrium path, there is a set of ω_s that offer Γ_2 with a strictly positive probability, and the buyer's posterior π is conditional on this set of ω_s. Also note that given any menu of contracts, it is characterized by the "on-the-equilibrium-path" belief of the seller, and the expected payoff of the seller is conditional on his true type. In order for any ω_s to not deviate from his equilibrium strategy, ω_s should find Γ_2 better than any Γ_2' he is not supposed to offer, in which case, the distribution of (ω_b, s) at the time he offers the second-period menu of contracts is the same conditional on the seller's true type. They might differ again after the buyer chooses a bundle. As for off-the-equilibrium-path Γ_2, any perfect Bayesian equilibrium needs to specify the buyer's belief on the seller type ω_s. This requires both equilibrium selection and further assumptions on Ω_s.

One can summarize the second period as Γ_2 that are offered with strictly positive probabilities, $\Omega_s(\Gamma_2)$ that offer Γ_2 with strictly positive probability, $\mu(\Gamma_2|\omega_s, \Gamma, \gamma)$ and $\pi(\omega, \omega_s|\omega_b, \Gamma, \gamma, s, \Gamma_2)$. Given these, the buyer's strategy $\sigma(\gamma_2|\omega_b, \Gamma, \gamma, s, \Gamma_2)$ maximizes his expected utility. From $V'' < 0$, there exists a unique optimal mechanism given the distribution over the buyer's posterior after observing Γ_2. However, Γ_2's need not be ex-post optimal given the buyer's posterior π's. And given that the seller's expected payoff is with respect to his true type, whenever a few types pool together, it is a knife-edge case that the buyer's posterior not knowing the precise true type will correspond to the seller's true type.

The second period by itself is closely related to static informed-principal problems, but Γ_2 really is the last step in signalling the seller's information, and the focus is the interaction between the two periods and Γ, γ, s being informative about ω.

The first period can also be summarized by Γ's offered with strictly positive probabilities and $\Omega_s(\Gamma)$ that offers each Γ with a strictly positive probability. The seller's strategy is $\mu(\Gamma|\omega_s)$, and the buyer's strategy is $\sigma(\gamma|\omega_b, \Gamma)$. As for the information structures, if it is conditional both on ω and ω_s, it allows the seller to specify the probabilities without revealing ω_s at the time. But the set of information structures allowed for the seller does matter for the results significantly. I mentioned in the model section that I restrict attention to partitioning Ω_s and adding martingales; the current set of results holds for a bigger class of information structures, but I do not list every case when they don't hold.

As for IC constraints in the first period, one can take care of deviations together in the same IC constraints. Both the seller's IC and the buyer's IC need to be satisfied, and from the seller's point of view, leaving aside detectable deviations for now, if the seller offers a menu of information structures that his type ω_s offers with zero probability on the equilibrium path, then the buyer's inference based on both Γ and the signal realization s are with respect to $\omega_s \in \Omega_s(\Gamma)$. The buyer's posterior in the beginning of the second period is the same as on the equilibrium path when the seller has a type in $\Omega_s(\Gamma)$. However, the seller offers Γ_2 in the second period, and if the seller offers Γ_2 that can be offered with a strictly positive probability after Γ, γ on the equilibrium path, then the buyer's inference is still the same as on the equilibrium path when $\omega_s \in \Omega_s(\Gamma) \cap \Omega_s(s|\gamma) \cap \Omega_s(\Gamma_2)$. This means that the distribution of (ω_b, s) ordered by $\mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma, s, \Gamma_2]$ is with respect to these ω_s's. However, the seller knows his true ω_s, and therefore, his expected utility is with respect to the distribution conditional on his true type. This shows that with undetectable deviations, the seller can change (i) the buyer's expectation on the seller's type ω_s, (ii) inference based on signal realization s, (iii) the posterior before choosing the bundle in the second period; in particular, (i) affects the buyer's participation decision and the choice of γ in the first period, and (ii),(iii) affect the participation and the consumption bundle in the second period. However, the seller is not changing the buyer's posterior alone. The buyer's posterior is on the joint distribution of ω, ω_s, and ω is the payoff relevant state. Since $\mathbb{E}[U(\omega)|\cdots]$ is the effective type in the second period, the seller is changing both the support and the distribution of the buyer type in the second period. Given any mechanism, one can compute the support and the distribution of $\mathbb{E}[U(\omega|\cdots)]$ for $(\Gamma, \gamma, \Gamma_2)$.

To take into account the buyer's incentives as well, in the second period, the buyer already knows $\mathbb{E}[U(\omega|\omega_b, \Gamma, \gamma, s, \Gamma_2)]$, and he chooses the bundle he would choose on the equilibrium path. This is because the buyer observes both ω_b, s privately, and Γ, γ, Γ_2 are offered publicly, so it really is the inference on ω, ω_s that's affected by the seller's deviation. Given that the buyer doesn't detect the seller's deviation, the buyer's inference is the same as on the equilibrium path, and therefore, his participation and the bundle choice are the same as on the equilibrium path. Therefore, the seller just needs to compare his expected revenue when his posterior is different knowing his own type.

Consider the buyer's incentives when the seller hasn't deviated in the first period. The buyer just maximizes his expected utility in the second period given his posterior after observing Γ_2. However, in the first period, given the set of Γ's, if he deviates in the first period and still participates, he will get a different information structure and a different piece of information from the signal. This leads to a different posterior of the buyer at the end of the first period, but now the seller will offer menus of contracts in the second period given Γ, γ and in particular assuming ω_b is one of buyer types who choose Γ, γ with a strictly positive probability on the equilibrium path. From the buyer's point of view, he just needs to update his posterior given the types of seller who offer each menu of contracts and maximizes his expected utility. Once the buyer has deviated, the seller assigns 0 probability on the buyer type and therefore, one cannot define what is one-stage deviation in this model. However, given that the buyer can still maximize his expected utility, as long as his participation decision doesn't change, the seller won't detect the deviation. But this shows that the usual notion of one-stage deviation principle in the mechanism design literature doesn't apply to this model, or more precisely, one cannot define what is "one-stage deviation."

For the rest of this section, I will derive properties of equilibria. I am not characterizing the set of all PBE, and as is often the case, not every equilibrium satisfies the same property. I will provide sufficient conditions for each property.

The following theorem shows that linearity of $U(\cdot)$ matters only if there is two-sided private information or interdependent values. Otherewise, one can normalize the parameters. The importance of this theorem is that when $U(\cdot)$ is linear, the effective buyer type in the second period is a martingale after purchasing an information structure in the first period, and therefore, if the seller is going to sell to all buyer types, there is no need to create more types and worsen the adverse selection problem.

However, one should also note that in most adverse selection problems satisfying Spence-Mirrlees condition, $U(\cdot)$ is rarely linear. Theorem 2 shows that when the first-period product is information, then one can normalize the second period as long as the seller has no private information and ω, q are multiplicatively separable.

Theorem 2. *In the buyer's second-period utility $U(\omega)V(q) - p$, the functional form $U(\cdot)$ implies the following:*

1. *When $U(\cdot)$ is linear and $\gamma : \Omega \to \Delta(S)$, the effective buyer type is always a martingale, i.e., $\mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma, s] = \mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma]$.*

2. *When $U(\cdot)$ is not linear, the effective buyer type is not a martingale for any generic π and γ, i.e.,$\mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma, s] \neq \mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma]$.*

3. *However, if the seller doesn't have any private information and $U(\cdot)$ is invertible, one can normalize $U(\omega)$ as ω and normalize π simultaneously. In this case, if $U(\cdot)$ is not invertible, there are payoff-equivalent states of the world that neither player can distinguish.*

4. *Therefore, any non-linear $U(\cdot)$ is only relevant with two-sided private information or correlated/interdependent values*

Theorem 2 doesn't show what happens if the seller doesn't sell to all buyer types he's facing, but theorem 7 in the following section shows that regardless of optimality, there are equilibria in which the seller can overcome the no-trade theorem as in Myerson-Satterthwaite (1983) by providing an information structure in the first period.

It also follows from limited commitment on the buyer side that if the buyer can choose whether to participate in each period and can choose to participate only in the second period without purchasing an information structure, the full-surplus extraction as in the information disclosure literature no longer holds.

Theorem 3. *When the buyer can choose whether to participate in each period, the seller can no longer extract full surplus.*

Since we are characterizing PBE, there are many inefficient equilibria that can be supported due to the buyer's belief after $\tilde{\Gamma} \notin \Pi(\Gamma), \tilde{\Gamma}_2 \notin \Pi(\Gamma_2)$. However, one can still compare the profit maximization in the usual monopoly setting without the seller's private information to the equilibrium of the second period. More specifically, there exists a unique optimal mechanism in the second period when $V'' < 0$, and if several seller types pool together in the second period, the posterior of each seller type doesn't have to be the same one that makes the menu of contracts optimal ex post. Ex-post inefficiency of mechanisms with full commitment has been shown in other papers as well, but in this case, ex-post inefficiency happens within a period due to the seller's private information.

Theorem 4. *Even though there exists a unique optimal mechanism without seller's private information, each menu of contracts Γ_2 offered in the second period need not be ex-post optimal in any equilibrium.*

The next theorem characterizes a sufficient condition for every equilibrium to have all types of seller pooling in the first period. This can be stated as a condition on equilibrium strategies, but there are sufficient conditions on primitives of the model that guarantee all equilibrium strategies satisfy this condition. When this condition is satisfied, every equilibrium must feature this property, and this can be considered as "implementation" in the mechanism design literature. One can also characterize conditions for all types of seller to pool in both periods in which case, from the buyer's perspective, he never learns anything about the seller's private information.

Theorem 5. *When the following conditions are satisfied, the only equilibrium strategy for the seller is pooling for all seller types:*
(i) each menu of information structures is a singleton, $\Gamma = \{\gamma\}$, (ii) (ω_b, s) ordered by $\mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma, s]$ in the second period is identical for all Γ, and (iii) the CDF of $\mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma, s]$ satisfies the first-order stochastic dominance with respect to Γ.

A special case is Bayesian persuasion ($\gamma : \Omega \to \Delta(S)$) and either the buyer has no private information or at least one of Ω_b, S is finite so that

$\mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma, s]$ *satisfies the first-order stochastic dominance.*

The first-order stochastic dominance and the ranking of information structures might seem like a strong condition, but this is satisfied with any information structure that is a mapping from Ω to $\Delta(S)$, e.g., "experiments" in Bayesian persuasion. Most papers in Bayesian persuasion don't assume that the designer has any private information himself, but in my model, it is sufficient that the information structure doesn't depend on ω_s. Whether the seller has any private information or not doesn't matter. This argument follows directly from the discussion of IC constraints before theorem 2.

Theorems 5 and 6 are the main results of this paper. One might argue the second period of my model is a fairly standard informed-principal problem except for the fact that the game changes if either player fully reveals their type in the first period. The fixed-point characterization of theorem 1 is different from other mechanism design problems, but the fixed point by itself doesn't show any property of the equilibria. Theorem 2 shows that there is a discrepancy between two-sided private information, or in the context of mechanism design, between the usual mechanism design problems and the informed-principal problems. And one can think of theorems 5 and 6 as "properties" of equilibria, as theorem 2 still shows the discrepancy but not the consequences of linearity.

The next theorem characterizes necessary conditions for creating more types to be optimal. One can also characterize necessary conditions for not creating endogenous types to be optimal and sufficient conditions for either case, creating or not creating endogenous types to be strictly better. However, given the assumptions of the model, it is not easy to characterize necessary and sufficient conditions for each. $\Omega, \Omega_s, \Omega_b$ are just assumed to be metric spaces, and π allows for any correlation across three objects. If one starts by assuming sufficient conditions for creating or not creating endogenous types to be optimal as modelling assumptions, then the theorems would be "necessary and sufficient" conditions, but this is really the matter of stating the same results.

Theorem 6. *Creating endogenous types, i.e., offering an information structure with multiple signals, is optimal only if the following conditions hold:*

1. *If $U(\cdot)$ is linear and $\gamma : \Omega \to \Delta(S)$, then the trade decision must change*
2. *If $U(\cdot)$ is linear, $\gamma : \Omega \times \Omega_s \to \Delta(S)$, and if the trade decision doesn't change, then the seller must signal his private information to increase the buyer's effective type*
3. *If $U(\cdot)$ is nonlinear and the trade decision doesn't change, then γ must increase the buyer's effective type; either $U(\cdot)$ is convex, or the seller must signal his private information.*

When the seller can signal his private information, given that the seller can commit to the probability of signal realization, one should consider the possibility of unraveling as in Milgrom (1981). In Milgrom (1981), the signals have a one-dimensional ordering, and the most favorable signal should be disclosed which leads to unraveling eventually. In my model, $\Omega, \Omega_s, \Omega_b$ are assumed to be metric spaces, and π can allow for arbitrary correlation across three variables; there need not be a one-dimensional ordering of the most favorable signal. However, if there is a seller type with the most favorable signal irrespective of the buyer type, then in points 2 and 3 of theorem 6, unraveling must happen.

3.3 Trade: $u_s(q|\omega) \neq 0$

Section 3.2 characterizes properties of PBE when $u_s(q|\omega) = 0$, i.e., the seller maximizes his revenue. When $u_s(q|\omega) \neq 0$, the seller's payoff from keeping the good himself depends on the state of the world. After the first period, even if the seller doesn't observe the signal realization that the buyer purchases, the seller learns about the buyer type from his choice of γ and therefore can update his belief about ω. In this case, the seller might not always want to trade with the buyer, which means that the limited-commitment solution will be different from the full-commitment solution except in the knife-edge case. Or to put it differently, compared to the previous section, the seller now faces one more constraint in any PBE, i.e., the seller's participation constraint in the second period. Also in characterizing the menus of information structures in the first period, the seller has incentives to learn from the experiment, which in the current model I do not allow the seller to elicit from the buyer directly in the second period. This would require an informed-principal version of the revelation principle which I am not aware of. Essentially, the seller's outside option in the second period depends on the state of the world, and the seller doesn't know the signal realization, but he knows the distribution of the signal realization; therefore, the seller can always offer the menu of contracts in the second period taking into account his posterior given the buyer types who choose each bundle in the equilibrium. This is another fixed point characterization, and uniqueness is not guaranteed.

When the seller's reservation value is non-zero, the difference comes from $u_s(q|\omega) \neq 0$ and the seller's participation decision in the second period. However, the seller's outside option, $\mathbb{E}[u_s(q|\omega)|\omega_s, \Gamma, \gamma]$, is related to papers on type-dependent outside options, see for example, Julien et al (2007) and Rochet (year). Furthermore, given that the buyer's participation decision and the choice of bundle in the second period signal his type ω_b further, the seller's outside option after the buyer's action in the second period will be again different from what it is given his information when he offers a menu of contracts. This type of learning has been shown more often in the common-value auction models.

Formally, the following theorem on Myerson-Satterthwaite no-trade theorm is the only characterization with $u_s(q|\omega) \neq 0$ in this paper. Complete characterization of pooling, separating, bunching or creating endogenous types are deferred to a different paper.

Theorem 7. *Suppose the valuations of the good of the seller and the buyer are the same as in Myerson-Satterthwaite (1983). When the seller can offer a menu of information structures in the first period, there are equilibria in which the seller and the buyer trade.*

4 Conclusion

I characterized properties of PBE in a two-period model where both the seller and the buyer have limited commitment and private information. The seller offers a menu of information structures in the first period and a menu of contracts in the second period. The buyer chooses an information structure in the first period and a bundle of price-quantity or price-quality in the second period; the buyer can also choose not to participate in either period. Since both players start with private information, this is a dynamic informed-principal problem, and given that an informed principal offers a menu of information structures in the first period, it is related to both information design and information disclosure; in general, putting a price tag on information hasn't been studied as extensively as with selling products, and information design and information disclosure literature don't charge for the information in most papers either.

Two-sided private information also allow for both correlated information, interdependent values as well as private values, and the key features of the model interact in non-trivial ways throughout.

First of all, the combination of informed principal and limited commitment makes it nontrivial to assume that Bester-Strausz (2001) will generalize immediately, i.e., the buyer might not report truthfully with every realized type, but it is without loss of generality to offer the set of types as the message space. It is further complicated by the fact that the first-period product matters for the second-period consumption utility from each bundle, and the complementarity is different again from the usual complementarity through $u(c_1, c_2)$. One can think about level effects and marginal utility, but strictly speaking, the first-period product, i.e., the information on the payoff-relevant state, changes both the support and the distribution of the effective buyer type, $\mathbb{E}[U(\omega)|\omega_b, \Gamma, \gamma, s, \Gamma_2]$, in the second period.

Properties of PBE show that two-sided private information endogenizes learning process differently from other learning models. Any experimentation paper with Poisson arrival and CARA utility or any dynamic moral hazard model with symmetric uncertainty on the payoff-relevant state assumes that the mechanism designer or the principal doesn't have any private information in the beginning of the game. The mechanism designer and the agent learn about the payoff-relevant state or the quality of the unknown arm by the agent exerting effort privately and both players observing the outcome. This means that the mechasnism designer's history coincides with the public history between the two players, and the agent's private history includes his own effort choices as well. On the other hand, most existing papers on the topic assume full commitment of the principal, and the payment for each outcome to the agent is specified in the beginning of the game. In my model, the seller and the buyer each have private information that are correlated with the payoff-relevant state, and they can learn from each other, in which case, if they can just put their private information together, that's the superset of the information available at the beginning of the first period. However, an information structure in the first period can offer "experiment" as in Bayesian persuasion literature as well, and the seller and the buyer can obtain a new piece of information that neither player knows in the beginning of the game. If both players had full commitment over both periods, one can ask whether there exists a mechanism that elicits both players' private information and also the signal realization truthfully. Then the question will be the split of the surplus from trade to maximize gains from trade. With limited commitment, there is another incentive to not disclose the private information fully in the first period. This by itself is not a new concept; it comes up with ratchet effect, and many models with fully persistent states often feature related issues. See for example, Rayo (2017) or Hörner-Skrzypacz (2016).

Now, in my model, there are further reasons that learning is more endogenous than existing papers. In addition to learning each other's private information, the buyer can learn more about the payoff-relevant state than the seller by observing the signal privately; however, the seller is the one who chooses the menu of information structures. One can even compare to delegation or information acquisition literature as well and compare different types to different preferences. The ex-post inefficiency as in theorem 4 has been shown in different contexts of deletaion or costly information acquisition literature.

References

[1] Bergemann, Dirk , Alessandro Bonatti, and Alex Smolin. (2018) "The Design and Price of Information," *forthcoming in American Economic Review.*

[2] Bergemann, Dirk, and Martin Pesendorfer. (2007) "Information Structures in Optimal Auctions," *Journal of Economic Theory*, 137: 580-609.

[3] Bester, Helmut, and Roland Strausz. (2001) "Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case," *Econometrica*, 69: 1077-1098.

[4] Che, Yeon-koo, and Johannes Hörner. (2018) "Optimal Design for Social Learning," *forthcoming in Quarterly Journal of Economics.*

[5] Condorelli, Daniele, and Balasz Szentes. (2018) "Information Design in the Hold-up Problem," *forthcoming in Journal of Political Economy.*

[6] Crawford, Vincent, and Joel Sobel. (1982) "Strategic Information Transmission," *Econometrica*, 1431-1451.

[7] Daskalakis, Constantinos, Alan, Deckelbaum, and Christos Tzamos. (2017) "Strong Duality for a Multiple-Good Monopolist," *Econometrica*, 85: 735-767.

[8] Eso, Peter, and Balazs Szentes. (2007) "Optimal Information Disclosure in Auctions and the Handicap Auction," *Review of Economic Studies*, 74: 705-731.

[9] Garicano, Luis, and Luis Rayo. (2017) "Relational Knowledge Transfer," *American Economic Review*, 107: 2695-2730.

[10] Garratt, Rod, and Marek Pycia. (2016) "Efficient Bilateral Trade," working paper.

[11] Gratton, Gabriele, Richard Holden, and Anton Kolotilin. (2017) "When to Drop a Bombshell," *forthcoming in Review of Economic Studies.*

[12] Julien, Bruno, Bernard Salanié, and Françoise Salanié. (2007) "Screening Risk-Averse Agents under Moral Hazard," *Economic Theory*, 151-191.

[13] Kamenica, Emir, and Matthew Gentzkow. (2011) "Bayesian Persuasion," *American Economic Review*, 101: 2590-2615.

[14] Krähmer, Daniel. (2018) "Full Surplus Extraction in Mechanism Design with Information Disclosure," working paper.

[15] Kremer, Ilan, Yishay Mansour, and Motty Perry. (2014) "Implementing the "Wisdom of the Crowd"," *Journal of Political Economy*, 122: 988-1012.

[16] Kojima, Fuhito, and Takuro Yamashita. (2017) "Double Auction with Interdependent Values: Incentives and Efficiency," *Theoretical Economics* 12: 1393-1438.

[17] Hart, Sergiu, and Phil Reny. (2010) "Revenue Maximization in Two Dimensions, or: Two Dimensions of Revenue Maximization," working paper.

[18] Hörner, Johannes, and Andy Skrzypacz. (2016) "Selling Information," *Journal of Political Economy*, 124: 1515-1562.

[19] Kwon, Suehyun. (2018) "Informed-Principal Problem in Mechanisms with Limited Commitment," working paper.

[20] Kwon, Suehyun. (2018) "Mechanism Design with Moral Hazard," working paper.

[21] Li, Hao and Xianwen Shi. (2017) "Discriminatory Information Disclosure," *forthcoming in American Economic Review.*

[22] Milgrom, Paul. (1981) "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics* 12: 380-391.

[23] Milgrom, Paul, and Robert Weber. (1982) "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50: 1089-1122.

[24] Milgrom, Paul, and Robert Weber. (1982) "The Value of Information in a Sealed Bid Auction," *Journal of Mathematical Economics*, 10: 105-114.

[25] Myerson, Roger B., and Mark A. Satterthwaite. (1983) "Efficient mechanisms for bilateral trading," *Journal of Economic Theory*, 29.2: 265-281.

[26] Mylovanov, Tymofiy, and Thomas Troeger. (2014) "Mechanism Design by an Informed Principal: The Quasi-Linear Private-Values Case," *Review of Economic Studies*, 81: 1668-1707.

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

[28] Rochet, Jean-Charles, and Lars Stole. (2002) "Nonlinear Pricing with Random Participation," *Review of Economic Studies*, 69: 117-146.

[29] Roesler, Anne-Katrin, and Balasz Szentes. (2017) "Buyer-Optimal Learning and Monopoly Pricing," *American Economic Review*, 107: 2072-2080.