Review of Economic Studies (2022) **0**, 1–41 https://doi.org/10.1093/restud/rdac070

Advance access publication 21 October 2022

Screening with Frames: Implementation in Extensive Form

FRANZ OSTRIZEK
University of Bonn
and
DENIS SHISHKIN
University of California San Diego

First version received April 2019*; Editorial decision July* 2022*; Accepted October* 2022 (*Eds.*)

We study a decision-framing design problem: a principal faces an agent with frame-dependent preferences and designs an extensive form with a frame at each stage. This allows the principal to circumvent incentive compatibility constraints by inducing dynamically inconsistent choices of the sophisticated agent. We show that a vector of contracts can be implemented if and only if it can be implemented using a canonical extensive form, which has a simple high–low–high structure using only three stages and the two highest frames, and employs unchosen decoy contracts to deter deviations. We then turn to the study of optimal contracts in the context of the classic monopolistic screening problem and establish the existence of a canonical optimal mechanism, even though our implementability result does not directly apply. In the presence of naive types, the principal can perfectly screen by cognitive type and extract full surplus from naifs.

Key words: Implementation, Screening, Framing, Extensive-form decision problems, Dynamic inconsistency, Sophistication, Naivete.

JEL Codes: D42, D82, D90, L12

1. INTRODUCTION

Ample evidence, casual empiricism, and introspection suggest that framing effects are common in choice. For example, decision makers tend to be risk averse in decisions framed as gains and risk seeking for losses (Tversky and Kahneman, 1981) and overestimate the impact of certain product attributes that are salient (Schkade and Kahneman, 1998).[1] In particular, the way a product is presented and the setting of the sales interaction—e.g. how the price is displayed and how much attention is focused on quality attributes of the product—can have a strong impact on

1. Also see the evidence discussed in Bordalo, Gennaioli and Shleifer (2013) and Kőszegi and Szeidl (2013), who provide models of consumer choice in which the properties of the choice set determine the attention on certain attributes.

The editor in charge of this paper was Andrea Galeotti.

consumer valuations.[2] Concordantly, many firms go to great expenses to improve the presentation of their product in largely non-informative and payoff-irrelevant ways through packaging, in-store design, and the emotions invoked by the sales pitch.

When preferences are affected by framing, a change of framing causes a change in preferences. Many economic interactions unfold in several stages admitting such changes. For example, when buying a car, a consumer is first exposed to a manufacturer's marketing material, contemplates his purchase decision at home, and is then affected by the way the product is presented by the dealer. Court trials involve plea-bargaining under the threat of higher charges potentially followed by court proceedings which may raise the defendant's hopes. Even a sales pitch itself unfolds sequentially. As a result, an agent's choices in such interactions with changing frames will display dynamic inconsistency. For instance, an alternative that is attractive under sales pressure may very well appear excessive when considered from the calm of one's home, prompting the agent to avoid the sales person entirely.

In many cases, the framing at different stages of the decision process and the resulting pattern of dynamic inconsistency is designed. In the examples above, this can be done by marketing teams, the legal system, and the sales person, respectively. In this article, we initiate the study of such decision-framing design: the structuring and framing of a decision problem to steer the choices of an agent with state- or frame-dependent preferences. We focus on a stylized single-agent setting in which the designer can commit and fully determine the sequence of frames and the contracts encountered by the agent. We characterize the implementable outcomes and apply our results to monopolistic screening. How does the designer use the power to affect tastes—only to more closely align the agent's tastes with her own or also to circumvent the incentive compatibility constraints usually implied by the agent's private information? Does this high level of influence granted to the designer imply that anything goes? What about the optimal pattern of frames? Does the designer—endowed with the opportunity to affect an agent's taste over an arbitrarily long and complex sequence of frames—require long and intricate patterns to eke out an advantage, or is there a simple canonical structure that carries all power of implementation?

We investigate these questions by adding framing and extensive forms to a principal-agent problem in a single-crossing quasi-linear environment. The principal chooses not only the contracts but also the structure of an extensive form along with a frame at each decision node. We assume that the agent is sophisticated, i.e. he correctly anticipates future choices, but chooses according to his current frame. By varying the frames, the mechanism can therefore create and exploit dynamic inconsistency. To illustrate how this allows the principal to bypass incentive compatibility constraints, consider to following example.

Example 1 *There are two types, θ^1 (low) and θ^2 (high), and two frames, ℓ (low) and h (high). Preferences of type θ^i in frame $f \in \{\ell, h\}$ are represented by $u^i_f(p,q) = \theta^i_f q - p$, where the marginal utility θ^i_f depends both on the type and the current frame (see Figure 1a).*

Consider the contracts $c^1 = (9,3)$ and $c^2 = (36,6)$ for types θ^1_l and θ^2_h, respectively. In particular, they correspond to the efficient full-extraction allocation that would arise if a principal with production cost $\frac{1}{2}q^2$ observed the types and was required to sell to the low type in the low

2. Consumer decisions are affected by the framing of insurance coverage (Johnson, Hershey, Meszaros and Kunreuther, 1993), the description of a surcharge (Hardisty, Johnson and Weber, 2010), whether discounts are presented in relative or absolute terms (DelVecchio, Krishnan and Smith, 2007), prices as totals or on a per-diem basis (Gourville, 1998), and by background music (Areni and Kim, 1993; North, Hargreaves and McKendrick, 1997; North, Shilcock and Hargreaves, 2003; North, Sheridan and Areni, 2016). Large effects of framing on consumer valuation are also found in incentivized lab experiments and across policy discontinuities (Bushong, King, Camerer and Rangel, 2010; Schmitz and Ziebarth, 2017).

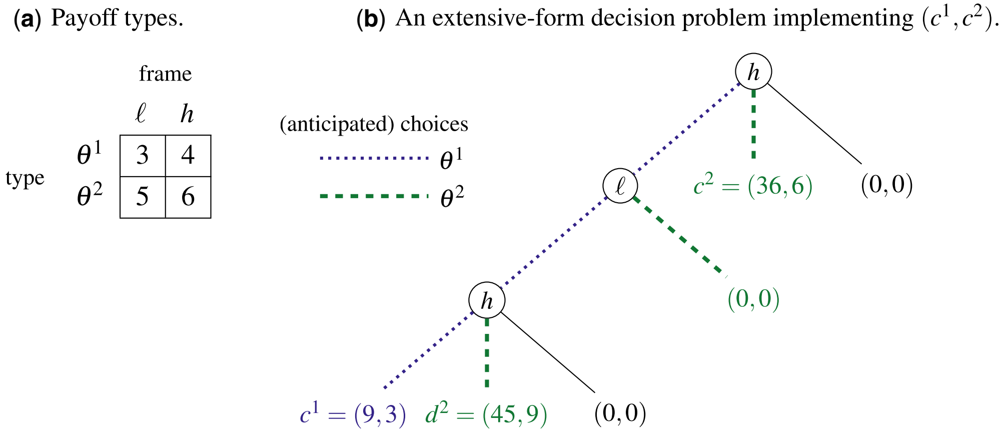

FIGURE 1
Example 1

frame and to the high type in the high frame. Clearly, this allocation is not incentive compatible for the high type in any fixed frame. It is however implementable in an extensive form that uses changes in framing: $h \to \ell \to h$.

To see how the principal achieves this, consider the extensive form in Figure 1b. It is easy to check that the low type prefers c^1 to any other contract in the extensive form in both frames and therefore proceeds through the tree to c^1. What about the high type? Because c^1 is preferable to c^2 for him in both frames, we need to show that such a deviation is infeasible *in this extensive form. To deviate to c^1, at the root the high type needs to choose the continuation problem leading to this contract. As he is sophisticated, he correctly anticipates his future choices but cannot commit. That is, at the second stage he anticipates that at the final stage he would pick the decoy d^2 (in the high frame). But according to his taste at the second stage (in the low frame), the decoy is unappealing, so he would choose the outside option. Hence, at the root the choice of the continuation problem is effectively equivalent to the outside option, thus, making the deviation to c^1 impossible.*

By placing a decoy contract as a "tempting poison pill" in the extensive form, the mechanism effectively removes the incentive compatibility constraint. This comes at the cost of adding an additional participation constraint, namely for the low type in the low frame, who has to pass through this frame on the path to his contract.

Generalizing the construction in the example, we identify a canonical extensive-form mechanism: For any finite number of types and frames, an allocation can be implemented if and only if it can be implemented in *three stages* using *two frames* (Theorem 1), provided that there is a sufficiently large set of feasible outcomes. Such canonical extensive forms have the following key features:

1. *Short interaction.* All types make at most three choices, and some only one.
2. *Fixed order of frames: high–low–high.* First, the agent is presented with a range of choices under the highest valuation frame. Some contracts are available immediately, some require a change of framing (to the second highest valuation frame) until a final decision is made back in the highest frame. The latter two stages involve a range of decoy contracts designed to throw off agents that misrepresented their type initially.

3. *Gains from framing: IR vs. IC.* The designer can either "reveal" (contract available at the root) or "conceal" (contract available in a continuation problem) each type. Therefore, she faces a trade-off between relaxing individual rationality by using only the highest frame (revealed types) and discouraging deviations by using frames in a high–low–high pattern to induce dynamic inconsistency (concealed types).

This final feature implies a direct characterization of implementability. For any vector of outcomes, there exist transfers such that it is implementable. In particular, outcomes do not have to satisfy the usual monotonicity requirement. Fixing transfers, a vector of contracts is implementable if and only if it satisfies a set of constraints. First, every contract needs to satisfy the participation constraint in the highest frame. Second, for every type one of two constraints has to be satisfied: either the participation constraint in the second highest frame, or that no other type—in the highest frame—is willing to imitate him.

Taking this single-agent multiple-self mechanism-design perspective contrasts with the growing literature analysing the impact of dynamic inconsistency in contracting when the pattern of taste changes is *given* by the agent's preferences (e.g. temptation or present bias): In our setting the pattern of inconsistency is *chosen* by the designer. We can hence analyse which patterns give the designer the most implementation possibilities and show that a simple pattern of taste changes (high–low–high) is sufficient.

We apply our results to the classic screening problem in Section 4. The principal designs a sales interaction to screen buyers by their taste for the product. Does she always present the product in the most favourable light or can changes of framing allow her to extract more surplus? While many factors influence the design of sales interactions, our application highlights the potential impact of decision-framing design. From a technical point of view, the space of feasible contracts in this application has a natural lower bound as the monopolist can only sell nonnegative quantities. Therefore, the space of outcomes is not sufficiently large to implement *every* vector of contracts in a canonical mechanism. We show that despite this restriction, the optimal allocation is implemented by a canonical extensive form (Theorem 2). To illustrate, consider again Example 1.

Example 1 (cont.) *Let us assume that the two types are equally likely and that costs are* $\frac{1}{2}q^2$. *If the principal offers a menu, this is a classic screening problem with an additional choice of a frame. It is easy to see that it is optimal to pick the high frame h and offer contracts so that* θ_h^1*'s participation constraint and* θ_h^2*'s incentive compatibility constraint bind, which yields a profit of* 20.[3]

The principal can do better. Using a canonical extensive form to circumvent the downward incentive compatibility constraint, the principal arrives at the mechanism described in Figure 1b. The high type does not obtain any information rent as the low type is concealed. Note that the decoy that achieves this has a non-negative quantity and is hence feasible. The maximal surplus that can be extracted from the low type is lower than in the static menu. There is a tradeoff between concealing the contract intended for the low type in the continuation problem and thereby eliminating information rents and extracting surplus from this type. In the present case, this reduction in total surplus is worth it for the principal, she obtains a profit of $22.5 > 20$.

In general, the profit-maximization problem is an optimization over the set of all extensive-form decision problems. Based on the structure of the optimal extensive form established in

3. In particular, the optimal contracts are $(p^1, q^1) = (8, 2)$ and $(p^2, q^2) = (32, 6)$. Note that with these functional forms, $q^i = \theta_f^i$ is efficient for frame f and the quality of type 2 is distorted downward compared to the efficient quality for *both* frames.

Theorem 2, we identify an equivalent optimization problem in contract space. The principal partitions the set of types into revealed and concealed. This partition determines the participation and incentive constraints: Concealing a type rules out the possibility of other types imitating it at the cost of a tighter participation constraint. In contrast to the classic setting, it is never optimal to exclude any type, as it is strictly better to sell a strictly positive quality to every type and conceal some of them instead (Proposition 3).[4]

For the main sections, we assume that agents are sophisticated. They correctly anticipate their choices but cannot commit to a course of action.[5] As the optimal sales interaction has a simple three-stage structure, correctly anticipating behaviour in this extensive form is relatively easy. Sophistication reflects the idea that consumers understand that they are more prone to choose a premium option when under pressure from the salesperson (high frame), and (in a low frame) avoid putting themselves in such situations that lead to excessive purchases. Moreover, consumers are exposed to sales pitches on a daily basis, they are experienced and understand the flow of the interaction. In addition, sophistication serves as a benchmark, by making it difficult for the principal to extract surplus. Even if consumers are fully strategically sophisticated and can opt out of the sales interaction at any point, framing in extensive forms affects the sales interaction and its outcomes. Indeed, the principal turns consumers' sophistication against them.

We also consider naive consumers (Section 4.4). They fail to anticipate that their tastes may change and choose a continuation problem as if their choice from this problem would be made according to their current tastes. For naive consumers, the principal can implement the efficient quantities in the highest frame and extract all surplus with a three-stage decision problem. She does so using decoy contracts in a bait-and-switch: Naive consumers expect to choose a decoy option tailored to them and reveal their type by choosing the continuation problem containing it at the root (bait), but end up signing a different contract due to the preference reversals induced by a change of frame (switch). When both naive and sophisticated consumers are present in arbitrary proportions and this cognitive type is not observable to the firm, our results generalize (Theorem 3).[6] The optimal extensive form still has three stages and implements the same contracts as if the cognitive type were observable. Neither sophisticated nor naive consumers gain information rents because of the presence of the other cognitive type.

Many jurisdictions mandate a right to return goods and cancel contracts, especially when the sale happened under pressure (e.g. door to door). This gives consumers the option to reconsider their purchase in a calm state of mind, unaffected by the immediate presence of the salesperson. In Section 4.5, we analyse such regulation and find that—while the principal can no longer use framing to exaggerate surplus—she can still use the resulting dynamic inconsistency to fully eliminate the information rent of all types. Sophisticated consumers do not require protection

4. This is in line with Corollary 2 in Salant and Siegel (2018), which states that there is no exclusion with two types, when the principal offers a framed menu under a participation constraint in a neutral frame. A related result is in Eliaz and Spiegler (2006). They show that there is no exclusion when the principal screens by the degree of sophistication. We show that no-exclusion holds when the principal screens by payoff type.

5. Another perspective is that a truly sophisticated agent foreseeing the changes of frame would instead become rational by taking an integrated, frame-independent point of view. While this is certainly possible in some cases, we focus on a purely procedural notion, which is consistent with the seminal work on time preference (Strotz, 1955; Laibson, 1997) and with the evidence that framing effects are observed within-subject (Tversky and Kahneman, 1981) and even among domain experts (Schwitzgebel and Cushman, 2015). Studying when and how agents "snap out of" their biases remains an important question for future research.

6. Spiegler (2011) notes that the principal can costlessly screen by cognitive type in a setting without taste heterogeneity. The underlying difference in the role of decoys—as poison pills for sophisticates and as bait for naifs—is parallel to the result that decoys cause sophisticates to act earlier and naifs to act later in task completion problems (Freeman, 2021).

by a right to return if they can decide to avoid the seller, e.g. by not visiting the store, but naive consumers would benefit even in this case.

1.1. *Related literature*

de Clippel (2014) studies Nash implementation with choice correspondences that cannot be derived from utility maximization but are not affected directly by the principal. Our setting is closer to multi-agent mechanism design with a multi-stage mechanism because of the presence of different frames. If we reinterpret our decision maker as a group of individuals with common knowledge of their type but different tastes, one individual corresponding to each frame, the principal applies implementation in backward induction without transfers. Herrero and Srivastava (1992) give abstract conditions for implementability in a general setting, we derive a canonical three-stage extensive form in the single-agent, multiple-self setting and derive properties of the optimal contracts in the classic screening problem.[7]

A growing literature studies the manipulation of framing by firms. Piccione and Spiegler (2012) and Spiegler (2014) focus on the impact of framing on the comparability of different products. Salant and Siegel (2018) study screening when framing affects the taste for quality, as in our setting. In this article, the principal chooses a framed menu, while we study the optimal design of an extensive-form decision problem to exploit the dynamic inconsistency generated by choice with frames and make predictions about the structure of interactions. In addition, our model makes different predictions for the use of framing and efficiency in the setting where the two are most closely comparable:[8] Using extensive forms, it is *always* more profitable to use framing (not only when it is sufficiently weak) and framing removes not only some but *all* distortions created by second-degree price discrimination in our setting.

Our article is also related to behavioural contract theory more generally (for a recent survey, see Kőszegi 2014, for a textbook treatment, see Spiegler 2011), in particular to screening problems with dynamically inconsistent agents (Eliaz and Spiegler, 2006, 2008; Esteban and Miyagawa, 2006a,b; Esteban, Miyagawa and Shum, 2007; Heidhues and Kőszegi, 2010, 2017; Zhang, 2012; Galperti, 2015; Moser and Olea de Souza e Silva, 2019; Yu, 2022).[9] These papers consider situations in which taste changes are *given* by the preferences of the agents (e.g. Gul and Pesendorfer (2001) or β-δ) and consequently design a two-stage decision problem as induced by the natural time structure of the problem.[10] We study how a principal *chooses*

7. Moore and Repullo (1988) show that subgame perfect implementation with multiple agents can be achieved using only three stages. This relies on monetary transfers between agents, however, which is not feasible between multiple selves. This restriction also turns the structure of the sequential problem on its head: the "test choice" happens on path while "challenging" the initial report terminates the interaction.

8. That is, comparing their Section 3 with our Section 4.5, where we impose a right to return the product in an exogenously given "neutral" frame. They also consider a model without returns but with a "basic" product that has to be offered and an insurance problem in which the monopolist can highlight one of the options, turning it into a reference point relative to which consumers experience regret.

9. Eliaz and Spiegler (2006, 2008) screen dynamically inconsistent agents by their degree of sophistication and optimistic agents by their degree of optimism, respectively. Esteban *et al.* (2007), Esteban and Miyagawa (2006a), and (Esteban and Miyagawa, 2006b) study screening when agents are tempted to over- or underconsume. Zhang (2012) studies screening by sophistication when consumption is habit inducing. Galperti (2015) studies screening in the provision of commitment contracts to agents with private information on their degree of time inconsistency, Heidhues and Kőszegi (2017) study selling credit contracts in this setting. Yu (2022) and Moser and Olea de Souza e Silva (2019) study optimal taxation problem, where agents are also heterogeneous in the degree of present bias.

10. For β-δ, two decision periods correspond to three periods including the final consumption period. Longer horizons are considered e.g. in Gottlieb and Zhang (2021), who show that in an insurance problem with symmetric information, all the inefficiency created by the exploitation of naivete is pushed to the final period. Following the pattern

the sequence of frames and an extensive form of arbitrary (finite) length to induce dynamic inconsistency and we show that a three-stage mechanism is optimal.

Given the optimal sequence of frames, this mechanism employs techniques similar to those in this literature. In particular, it involves off-path options that remain unchosen by every type ("decoys"). In Esteban and Miyagawa (2006a) and Galperti (2015) such decoys make deviations less attractive and are thus analogous to the decoy contracts, we introduce in the mechanism for sophisticated agents. Heidhues and Kőszegi (2010) show that credit contracts for partially sophisticated quasi-hyperbolic discounters feature costly delay of the payment which the consumer fails to expect when signing the credit contract. Immediate repayment is hence an unused option analogous to the "bait" decoys we introduce to screen naive consumers.

Glazer and Rubinstein (2012, 2014) consider models where the principal designs a procedure such that misrepresenting their type is beyond the boundedly rational agents' capabilities. While their decision problems are based on hypothetical questions about the agent's type, we show that it is possible to structure a choice problem with framing to make it impossible to imitate certain types.

There is a large literature on endogenous context effects, e.g. through focusing the attention of the decision maker on attributes that vary strongly or are exceptional within the choice set (Bordalo *et al.*, 2013; Kőszegi and Szeidl, 2013). We consider the case of framing through features of the choice situation, such as the sales pitch or the presentation format. Thus, consumers in our model fit into the choice with frames framework of Salant and Rubinstein (2008).

Dynamic sales interactions can also be analysed from the perspective of information provision (e.g. Eső and Szentes, 2007; Li and Shi, 2017; Wei and Green, 2022). Information needs to satisfy a martingale condition and consumers remain dynamically consistent, while we focus on the dynamic role of framing to relax incentive compatibility by inducing dynamic inconsistency.

2. FRAMES AND EXTENSIVE FORMS

This section introduces our quasi-linear single-agent mechanism design framework with extensive forms and frames.

2.1. *Contracts and frames*

A contract c is a pair of a transfer $p \in \mathbb{R}$ and an outcome q from an interval $\mathcal{Q} \subseteq \mathbb{R}$. The space of contracts is $\mathcal{C} = \mathbb{R} \times \mathcal{Q}$. Anticipating the application to monopolistic screening, in discussions we sometimes refer to the principal as the seller, the agent as the consumer, q as quality and p as price.

There is a finite set of frames F with $|F| \geqslant 2$ and a finite type space Θ with $|\Theta| \geqslant 2$. Each type is a function $\theta: F \to \mathbb{R}_{++}$ that maps frames into payoff types, denoted as $\theta_f := \theta(f)$. We assume that frames are distinct in the sense that no two frames induce the same vectors of payoff types. For a given payoff type θ_f, the consumer is maximizing a quasi-linear utility function

$$u_{\theta_f}(p,q) := v_{\theta_f}(q) - p,$$

of preference reversals ingrained in the β-δ specification, the allocation cascades, always postponing the consumption shortfall by one period. This allows virtual efficiency in consumption smoothing problems as the number of periods grows, but creates unrectifiable inefficiency in an effort choice problem. We show that for a single decision problem with (partially) naive agents, the optimal taste change pattern is h–ℓ–h, which allows the designer to achieve full surplus extraction despite private information.

where $v\colon \mathbb{R}_{++}\times\mathcal{Q}\to\mathbb{R}$ is a twice differentiable function that is strictly increasing and concave in q. We also assume that the environment satisfies a version of a single-crossing property. In particular, we assume that v has increasing marginal differences that are bounded away from zero, that is, there exists $\varepsilon>0$ such that $\frac{\partial^2 v}{\partial\theta_f\partial q}\geqslant\varepsilon$. For example, $v_{\theta_f}(q)=\theta_f q$ is a widely used functional form satisfying these assumptions. Note that this model of framing can accommodate frames affecting price perception, as they change the subjective tradeoff between outcomes and transfers.

The agent has an outside option which we normalize to $\mathbf{0}:=(0,0)$ and we assume $0\in\mathcal{Q}$. We also normalize $v_{\theta_f}(0)=0$, for all payoff types θ_f.[11] We say a vector of contracts $\mathbf{c}=(c_\theta)_{\theta\in\Theta}=\big((p_\theta,q_\theta)\big)_{\theta\in\Theta}$ is *non-negative* and write $\mathbf{c}\geqslant 0$ if all outcomes q_θ are non-negative. We refer to the constraints

$$u_{\theta_f}(c_\theta)\geqslant 0,\ \text{and} \tag{P^f_θ}$$

$$u_{\theta_f}(c_\theta)\geqslant u_{\theta_f}(c_{\theta'}) \tag{$\mathrm{IC}^f_{\theta\theta'}$}$$

as the participation constraint for θ and the incentive compatibility constraint from θ to θ' in frame f, respectively.

Our central assumption is that the frames and types are ordered.

Assumption 1 (Comonotonic environment) *For any types $\theta,\theta'\in\Theta$ and frames $f,f'\in F$,*

$$\theta_f>\theta_{f'}\implies\theta'_f>\theta'_{f'}\quad \textit{and}\quad \theta_f>\theta'_f\implies\theta_{f'}>\theta'_{f'}.$$

The first part of the assumption implies that frames can be ordered by their impact on the valuation. There is a lowest frame, i.e. a frame inducing the lowest valuation for every type and a highest frame, i.e. a frame inducing the highest valuation for every type. The second part implies that types can also be ordered by their valuation independently of the frame. With slight abuse of notation, we denote the order on frames and types using regular inequality signs.

In many cases, a frame has a similar impact on different consumer types. The more effectively a seller emphasizes quality, for instance, the higher a consumer values quality irrespective of their type. The first part of our assumption is satisfied as long as the *direction* of the impact of a given frame is the same for all types. The second part is satisfied as long as the *size* of the effect is not too different between types relative to their initial difference in valuation.

Assumption 1 precludes any frame from impacting the valuations of different types in a different direction. For example, focusing a car buyers attention on emissions may increase the valuation of a "green" car for some buyers while reducing the valuation of all cars, including the "green" car, for others. Similarly, it rules out cases where the order of types by their payoff parameter depends on the frame. For example, the demand for health insurance coverage may be lower among smokers than non-smokers if they are not reminded about the long run effects of their habit, but is higher for smokers than non-smokers if the effects of smoking are made salient during the sale of insurance. It also rules out that certain frames are specific to certain types. We discuss how we can relax our assumptions in Section 3.4.

We now illustrate with two examples how our model encompasses framing effects encountered in applications. First, consider a model of attention to attributes. The consumers' attention

11. Starting with an outside option (q_0,p_0) and an unnormalized $\tilde{v}$, we can always set $v(q,\theta_f):=\tilde{v}(q_0+q,\theta_f)-\tilde{v}(q_0,\theta_f)+p_0$. This transformation preserves our assumptions. Note that this normalization is without loss as we assume an outside *option* (instead of an outside utility level) and that this outside option is fixed (not type or frame dependent).

can be directed towards the quality of the products (f_q), towards the price (f_p), or be neutral (f_n), induced e.g. through the way the options are presented by the salesperson, the physical presence of the object (Bushong *et al.*, 2010), or the price format (Gourville, 1998; Schmitz and Ziebarth, 2017). Let $\lambda > 1$ parametrize the impact of attention. The consumer evaluates his options according to

$$u_{\theta_f}(q,p) = \begin{cases} \lambda\theta v(q) - p & \text{if } f = f_q \\ \theta v(q) - p & \text{if } f = f_n \\ \theta v(q) - \lambda p & \text{if } f = f_p \end{cases}$$

for an increasing and concave function v. It is easy to see that we can find a suitable representation by writing $\theta_{f_q} = \lambda\theta$, $\theta_{f_n} = \theta$, $\theta_{f_p} = \frac{\theta}{\lambda}$, and $u_{\theta_f}(q,p) = \theta_f v(q) - p$.

As a second example, consider the sale of insurance. The type θ denotes the probability of the damage of fixed size ℓ. The agent has background wealth of w and is subject to loss aversion. Given a reference point r, he evaluates an insurance coverage of q at price p according to

$$u_{\theta_f}(q,p) = \theta V(w - L + q - p, r) + (1-\theta)V(w - p, r).$$

Suppose there are two possible frames, which determine the reference point. A loss frame (ℓ) inducing $r \in (w - L + q - p, w - p)$, i.e. a higher insurance coverage is seen as reducing the loss and a gain frame (g) inducing $r = w - L$, i.e. the agent has internalized the loss and the insurance coverage is seen as a gain. This framing corresponds to presenting policies in terms of deductibles (loss) or rebates (gains) in Johnson *et al.* (1993).[12] Using the common piece-wise linear specification for the value function, and imposing a domain restrictions on the contracts such that an $r \in (w - L + q - p, w - p)$ exists, these preferences are represented by $u_{\theta_f}(q,p) = \theta_f q - p$ with $\theta_g = \theta$ and $\theta_\ell = \frac{\lambda\theta}{1+\theta(\lambda-1)}$.

2.2. *Extensive-form decision problems*

A mechanism is an extensive-form decision problem (EDP) with a frame attached to each decision node. For example, the following situation can be represented by a two-stage extensive-form decision problem. First, the consumer contemplates whether to visit the store and then purchases a product in the store. Perhaps, the consumer is initially affected by marketing materials (frame at the root) and then the consumer is affected by the sales pitch in the store (frame at the second stage).

Formally, an EDP is a perfect- and complete-information extensive-form game with perfect recall where the players are the multiple selves of the agent corresponding to different frames and the outcomes are contracts. We provide a formal definition below to introduce the notation used in the analysis. Let $\mathcal{E}^k$ denote the set of all k-EDPs, i.e. extensive-form decision problems with up to k stages. For notational convenience, let $\mathcal{E}^0 := \mathcal{C}$. For any set S, let $\mathcal{P}(S)$ denote the set of all finite subsets of S containing the outside option $\mathbf{0}$. Then construct the set of 1-EDPs as

$$\mathcal{E}^1 := \mathcal{P}(\mathcal{E}^0) \times F,$$

12. See also Gottlieb and Mitchell (2020), who show that the susceptibility to narrow framing (corresponding to a state-dependent reference point such that individuals perceive the insurance premium as a loss and the net insurance payout as a gain) is an important determinant of purchasing long-term care insurance, dwarfing the effect of risk aversion and adverse selection.

that is, a 1-EDP $e=(A,f)$ is a pair of a finite menu of contracts A and a frame f. For each $k>1$, the set of k-EDPs is

$$\mathcal{E}^k := \mathcal{P}\big(\cup_{l=0}^{k-1}\mathcal{E}^l\big)\times F,$$

so that a k-EDP $e=(E,f)$ is a pair of a finite set E of EDPs with less than k stages and a frame f.[13] Finally, the set $\mathcal{E}$ of all finite EDPs is given by

$$\mathcal{E} := \bigcup_{k=1}^{\infty}\mathcal{E}^k.$$

To illustrate, in terms of our notation, the EDP in Example 1 is $\big(\big\{\big(\big\{\big(\big\{c^1,d^2,\mathbf{0}\big\},h\big),\mathbf{0}\big\},\ell\big),\mathbf{c}^2,\mathbf{0}\big\},h\big)$.

Finally, note that we require the outside option to be available at each stage. In other words, the agent can end the interaction at any point. An extensive-form mechanism models a single binding decision that the agent arrives at in several steps, e.g. the interaction with an insurance agent leading to the signing of the contract. In such situations, forcing the agent to continue to participate in the mechanism risks backlash and is also legally challenging as pre-contractual duties are limited and damages are speculative. Methodologically, with such a strong form of participation constraint, our model makes it harder for the principal to use framing effects for implementation, and our results can be viewed as providing a lower bound on the set of implementable contracts and maximal profits. We discuss the outcome with alternative assumptions on participation in Sections 3.4 and 4.6 and analyse the case with a participation constraint in a fixed frame in Section 4.5.

2.3. *Choice from extensive-form decision problems*

We now introduce our main solution concept. It is immediate to define the choices of an agent of type θ in a 1-EDP (A,f) as the θ_f-optimal contracts in A. To define the agent's choice for all EDPs, we need to make an assumption on how he anticipates his choices at subsequent decision nodes. For the main part of the article, we assume that the agent is *sophisticated*.[14] That is, presented with a choice between several continuation problems, the agent correctly anticipates future choices, but chooses the continuation problem according to her current frame. The current self has no commitment power other than the choice of a suitable continuation problem at the current stage.

Formally, when an EDP e is seen as a game between multiple selves, the set $\Sigma^e\subseteq\mathcal{C}^\Theta$ of *solutions* of e is the set of vectors of subgame perfect Nash equilibrium outcomes, i.e. vectors consisting of contracts obtained via backward induction for each type. The definition is standard, but we provide one for completeness.

First, a mapping $\sigma:\ \Theta\to\mathcal{C}$ is a *solution* of a 1-EDP $e=(A,f)\in\mathcal{E}^1$, i.e. $\sigma\in\Sigma^e$, if and only if $\sigma(\theta)$ maximizes u_{θ_f} on A for all θ. Now take any $k>1$ and suppose that Σ^e is well-defined for all $e\in\cup_{l=0}^{k-1}\mathcal{E}^l$.[15] Then consider a consumer facing a k-EDP $e_k=(E,f)\in\mathcal{E}^k$. Choosing between continuation problems in E, she anticipates her choice $\sigma^e(\theta)\in\mathcal{C}$ in each $e\in E$, but evaluates the contracts $\{\sigma^e(\theta)\}_{e\in E}$ in the current frame f. Thus, σ is a solution of e_k if there exists a solution

13. Note that the choice of frame at each stage is unrestricted. In particular, the frame at the root of an EDP does not place any constraints on subsequent frames. We discuss the role of this assumption in Section 3.4.

14. We analyse naive consumers in the context of our screening application in Section 4.4.

15. For notational convenience, let $\Sigma^{e_0}=\{\theta\mapsto e_0\}, \forall e_0\in\mathcal{E}^0=\mathcal{C}$.

$\sigma^e \in \Sigma^e$ for every $e \in E$, such that for all $\theta \in \Theta$

$$\sigma(\theta) \in \underset{e \in E}{\operatorname{argmax}} \, u_{\theta_f}(\sigma^e(\theta)).$$

We say that a vector of contracts $\mathbf{c}$ is *implemented* by an EDP e if $\mathbf{c}$ is a solution of e. We call $\mathbf{c}$ *implementable* if it is implemented by some EDP.

3. IMPLEMENTABLE CONTRACTS

In this section, we provide a characterization of all implementable contracts. Before analysing the general problem, it is instructive to consider two special cases.

3.1. *Two trivial special cases*

First, consider a "*single-stage*" setting, in which the principal can use only choose a 1-EDP. Clearly, in this case the principal cannot use multiple frames. Second, consider a "*single-frame*" setting, in which the EDP must use the same frame at every stage. The extensive-form structure does not matter in this case: As the agent is perfectly rational and dynamically consistent, she picks the most preferred contract available in the extensive form. Hence, an extensive form is equivalent to an unstructured menu offering the same set of contracts.

In both cases, the implementation problem is a standard static problem. That is, a revelation principle applies, and implementability is equivalent to incentive compatibility and participation constraints being satisfied in some frame.

Observation 1 *For each $\mathbf{c} \in \mathcal{C}$, the following are equivalent:*

(i) $\mathbf{c}$ *is implementable by a* 1*-EDP.*
(ii) $\mathbf{c}$ *is implementable by an EDP using only a single frame.*
(iii) $\mathbf{c}$ *satisfies the* $\text{IC}^f_{\theta\theta'}$ *and* P^f_θ *constraints for all* θ, θ' *for some* $f \in F$.

We provide all omitted proofs in Appendix A.

Individually, framing and extensive forms do not qualitatively affect the implementation possibilities. It is only through their interaction that they realize their potential.

3.2. *Canonical extensive forms*

In this section, we show that despite the complexity of the environment, contracts can be implemented using a simple three-stage structure. Towards this result, we define a class of EDPs which share structural features. Define the *high* and *low* frames h and ℓ as the highest and second highest frames

$$h := \underset{F}{\operatorname{argmax}} \, \theta, \quad \ell := \underset{F \setminus \{h\}}{\operatorname{argmax}} \, \theta,$$

for some $\theta \in \Theta$, respectively, and note that these definitions are independent of θ under Assumption 1.

First, towards a definition of canonical EDPs, partition the set of types Θ into two sets corresponding to the two ways to present the contract associated to a given type: Contracts c_θ for *revealed types* ($\theta \in \Theta_R$) are presented at the root, while contracts for *concealed types* ($\theta \in \Theta_C$) are presented in separate continuation problems e_θ. Then, the three stages are (see Figure 2):

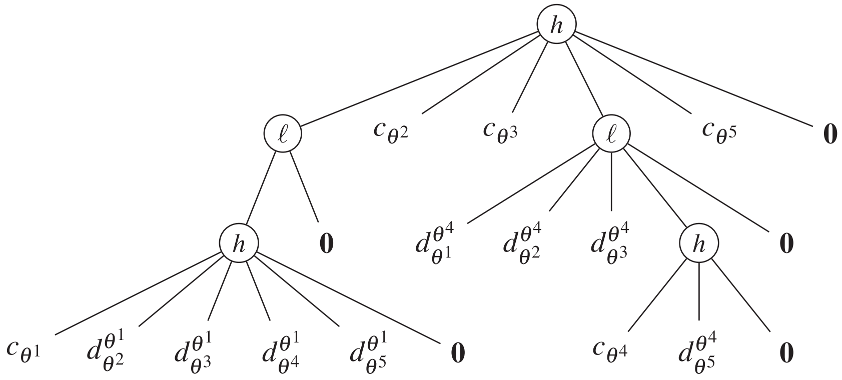

FIGURE 2

A canonical EDP for $(c_{\theta^1},\dots,c_{\theta^5})$ with $\Theta_R=\{\theta^2,\theta^3,\theta^5\}$ and $\Theta_C=\{\theta^1,\theta^4\}$.

1. Root: choose in the high frame between contracts for revealed types, continuation EDPs for concealed types, and the outside option.
2. Continuation problem for a concealed type θ: choose in the low frame between decoys for types below θ, continue to the terminal choice for θ, and the outside option.
3. Terminal choice for a concealed type θ: choose in the high frame between a contract for θ, decoys for types above θ, and the outside option.

Definition 1 *An EDP e is a* canonical *EDP for a vector of contracts* $\mathbf{c}$ *if there exists a partition* $\{\Theta_C,\Theta_R\}$ *of* Θ*, and decoy contracts* $\{d^{\theta}_{\theta'}\}_{\theta\in\Theta_C,\theta'\neq\theta}$*, such that*

$$e=\left(\{e_\theta\}_{\theta\in\Theta_C}\cup\{c_\theta\}_{\theta\in\Theta_R}\cup\{\mathbf{0}\},h\right),\ \textit{where} \tag{1}$$

$$e_\theta=\left(\left\{\left[\{c_\theta,\mathbf{0}\}\cup\{d^{\theta}_{\theta'}\}_{\theta'>\theta},h\right],\{\mathbf{0}\}\cup\{d^{\theta}_{\theta'}\}_{\theta'<\theta}\right\},\ell\right),\forall\theta\in\Theta_C. \tag{2}$$

The extensive form in Example 1 is a canonical EDP. Type θ^1 is concealed—his contract is available only after a continuation problem—while type θ^2 is revealed—his contract is available immediately at the root.

Our first main result provides a characterization of the implementable vectors of contracts. We say that a statement holds if the space of outcomes is sufficiently large if there exists a finite interval $[\underline{q},\overline{q}]$ such that the statement holds for any $\mathcal{Q}\supseteq[\underline{q},\overline{q}]$.

Theorem 1 *For each non-negative vector of contracts* $\mathbf{c}$*, if the space of outcomes is sufficiently large, then the following are equivalent*

(i) $\mathbf{c}$ *is implementable,*
(ii) $\mathbf{c}$ *is implementable by a canonical EDP,*
(iii) $\mathbf{c}$ *satisfies the constraints* $\{\mathrm{P}^h_\theta\}_{\theta\in\Theta_R}$, $\{\mathrm{P}^\ell_\theta\}_{\theta\in\Theta_C}$*, and* $\{\mathrm{IC}^h_{\theta\theta'}\}_{\theta\in\Theta,\theta'\in\Theta_R}$ *for some partition* $\{\Theta_C,\Theta_R\}$ *of* Θ.

Proof. Statement (ii) trivially implies (i) for any space of outcomes. It is, therefore, sufficient to show that (i) implies (iii) and (iii) implies (ii) for some a sufficiently large quality domain. We establish these implications in Propositions 1 and 2 below. □

This result implies that the principal can always use an EDP with a simple structure. First, implementation can be achieved in *three stages for an arbitrary number of agent types*, even though the principal has arbitrarily complex and long extensive forms at her disposal. As the number of types increases, the structure and length of the decision problem stays the same, only the number of available contracts increases.

Second, types are *separated at the root*. The principal does not use the extensive-form structure to discover the type of an agent piecemeal, it is an implementation device to protect contracts against imitation.

Third, only the *two highest frames* are used. As we have seen in Observation 1, if every decision node uses the same frame, the extensive-form structure is irrelevant for the agent's choice. Consequently, the principal uses at least two frames in order to induce violations of dynamic consistency. As long as the principal induces such violations, the decoys can be constructed irrespective of the number of or cardinal differences between the frames used. Hence, two frames are sufficient for the principal to reap *all potential gains* from such violations. Finally, only the highest two are used in the optimal EDP in order to relax the participation constraints.

Finally, implementability does not put any restriction on the vector of outcomes. This is in contrast to the classic setting with dynamically consistent agents where implementability typically implies monotonicity.

Corollary 1 For every non-negative vector of outcomes $\mathbf{q}$, if the space of outcomes is sufficiently large, then there exists a vector of transfers $\mathbf{p}$ such that $(\mathbf{p}, \mathbf{q})$ is implementable.

3.3. *Necessary and sufficient conditions for implementation*

In order to provide the foundation for Theorem 1, we proceed in two steps. First, we identify necessary conditions that every implementable vector of contracts has to satisfy. Then, we show that the necessary conditions are sufficient to ensure that the contract can be implemented by a canonical EDP. In particular, we explicitly construct decoy contracts and show that the principal can thereby eliminate all IC constraints into concealed types.

3.3.1. Necessary conditions for implementation by general EDPs. Consider an arbitrary EDP implementing a vector of contracts $\mathbf{c} = (c_\theta)_{\theta \in \Theta}$. Denote the frame at the root by f_R. Extending the notion of revealed and concealed types from canonical EDPs, for each type θ there are two possibilities: If there exists a path from the root to c_θ with all decision nodes set in f_R, then θ is called revealed. Alternatively, if every path from the root to c_θ involves at least one $f_\theta \neq f_R$, then θ is called concealed. As usual, we will denote the sets of revealed and concealed types by Θ_R and Θ_C, respectively.

First, consider the participation constraints. If the path from the root to c_θ passes through a node in frame f, then, since the outside option is always available, c_θ needs to satisfy the corresponding participation constraint P_θ^f. In particular, every contract has to satisfy the constraint at the root $P_\theta^{f_R}$.

We now turn to the incentive compatibility constraints. If θ is revealed, c_θ can be reached by any type from the root, as consumers are dynamically consistent when the frame does not change along the path. Consequently, for any θ', c_θ must not be an attractive deviation, that is

$$u_{\theta'_{f_R}}(c_{\theta'}) \geqslant u_{\theta'_{f_R}}(c_\theta), \qquad (\mathrm{IC}_{\theta\theta'}^{f_R})$$

for all $\theta' \in \Theta$ and $\theta \in \Theta_R$.

If θ is concealed, there is a change of frame along the path to c_θ. This induces a violation of dynamic consistency, which may make deviations into c_θ impossible. As we are looking for necessary conditions, we impose no incoming IC constraint in this case.

The argument so far identifies a family of conditions indexed by $(f_R, \{f_\theta\}_{\theta\in\Theta}, \Theta_C)$, such that a vector of contracts is implementable only if it satisfies at least one member of the family. The following proposition shows that for non-negative $\mathbf{c}$, we can always set $f_R = h$ and $f_\theta = \ell$ for a suitably chosen Θ_C. First, we can set $f_R = h$ by concealing all types if necessary. Then, the exact frames $\{f_\theta\}_{\theta\in\Theta}$ only affect the participation constraints, which are relaxed by moving to higher frames, while the change of frame bypasses IC. Consequently, the contract must satisfy the least restrictive participation constraints, i.e. in the highest and second highest frame.

Proposition 1 *If a non-negative vector of contracts* $\mathbf{c}$ *is implemented by an EDP, then it satisfies the constraints* $\{\mathrm{P}^h_\theta\}_{\theta\in\Theta_R}$, $\{\mathrm{P}^\ell_\theta\}_{\theta\in\Theta_C}$, *and* $\{\mathrm{IC}^h_{\theta\theta'}\}_{\theta\in\Theta,\theta'\in\Theta_R}$ *for some partition* $\{\Theta_C, \Theta_R\}$ *of* Θ.

The necessary conditions illustrate the trade-off between using framing to relax individual rationality and its use to discourage deviations. For revealed types, the participation constraint needs to be satisfied only in the highest frame, the frame resulting in the least restrictive constraint. This results in the largest set of individually rational contracts. For concealed types, the participation constraint needs to be satisfied in the second highest frame. This shrinks the set of individually rational contracts. The principal is compensated for this reduction through the removal of IC constraints into concealed types.

3.3.2. Sufficient conditions for implementation by canonical EDPs. To construct a canonical EDP that implements a vector of contracts $\mathbf{c}$, we proceed in two steps. First, we need to determine the set of concealed types. Clearly, type θ can be concealed in a canonical EDP only if c_θ satisfies the participation constraint P^ℓ_θ, since otherwise he would prefer to opt out in the second stage. Second, for each concealed type θ, we construct a continuation problem using decoys which make *all* deviations intro c_θ impossible.

Proposition 2 *If a non-negative vector of contracts* $\mathbf{c}$ *satisfies the constraints* $\{\mathrm{P}^h_\theta\}_{\theta\in\Theta_R}$, $\{\mathrm{P}^\ell_\theta\}_{\theta\in\Theta_C}$, *and* $\{\mathrm{IC}^h_{\theta\theta'}\}_{\theta\in\Theta,\theta'\in\Theta_R}$ *for some partition* $\{\Theta_C, \Theta_R\}$ *of* Θ, *then* $\mathbf{c}$ *is implementable by a canonical EDP if* $\mathcal{Q}$ *is sufficiently large.*

As in Example 1, the principal constructs decoys for the final stage to render downward deviations into concealed types impossible in the extensive form. She furthermore constructs decoys for the intermediate stage to render upward deviations into concealed types impossible as well. These constructions are the central step in our results, and we therefore present them in the text. They ensure that if θ is concealed, no type $\theta' \neq \theta$ can successfully imitate θ. Neither c_θ nor the decoys in e_θ interfere with the choices of any other type as they are dominated by the outside option at the root (i). Furthermore, θ chooses the intended contract (ii). This construction requires a sufficiently large space of outcomes.

Lemma 1 (Decoy construction) For any $\theta\in\Theta$ and non-negative $c_\theta \in \mathcal{C}$, c_θ satisfies P^ℓ_θ if and only if there exist decoys $(d^\theta_{\theta'})_{\theta'\neq\theta}$, such that the corresponding e_θ in (2) has a solution σ that satisfies

(i) $\sigma(\theta') \preccurlyeq_{\theta_h} \mathbf{0}$ for all $\theta' \neq \theta$, and
(ii) $\sigma(\theta) = c_\theta$.

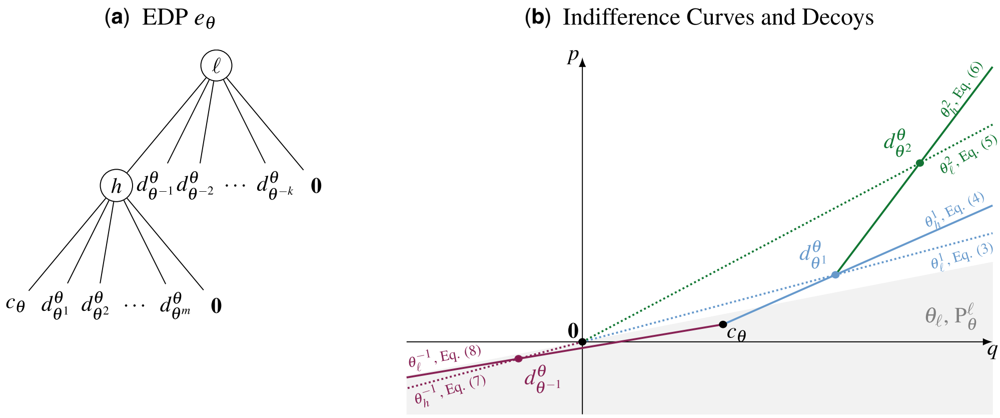

FIGURE 3
The construction of e_θ.

Proof. Construction. The construction of the decoys and the continuation problem e_θ is illustrated in Figure 3. At the terminal stage, agents are presented with the choice between the contract c_θ, the outside option and a set of decoys $\{d^\theta_{\theta'}\}_{\theta'>\theta}$, one for every type greater than θ. Given a contract c_θ, the decoy $d^\theta_{\theta^1}$ for the next largest type θ^1 has to satisfy

$$u_{\theta^1_\ell}(\mathbf{0}) \geqslant u_{\theta^1_\ell}(d^\theta_{\theta^1}). \tag{3}$$

$$u_{\theta^1_h}(c_\theta) \leqslant u_{\theta^1_h}(d^\theta_{\theta^1}) \tag{4}$$

For a parsimonious construction, we pick $d^\theta_{\theta^1}$ at the intersection of the two indifference curves (Figure 3b). Then, the decoy $d^\theta_{\theta^2}$ for the next type θ^2 solves

$$u_{\theta^2_\ell}(\mathbf{0}) = u_{\theta^2_\ell}(d^\theta_{\theta^2}). \tag{5}$$

$$u_{\theta^2_h}(d^\theta_{\theta^1}) = u_{\theta^2_h}(d^\theta_{\theta^2}) \tag{6}$$

Proceeding by induction, we construct decoys for all $\theta' > \theta$.

At the root of e_θ, agents are presented with the choice between the continuation, the outside option, and a set of decoys $\{d^\theta_{\theta'}\}_{\theta'<\theta}$, one for every type smaller than θ. Similarly to the above, we now proceed downwards from c_θ. The decoy $d^\theta_{\theta^{-1}}$ for the next smaller type θ^{-1} is implicitly defined by the system

$$\max\{u_{\theta^{-1}_\ell}(c_\theta), 0\} = u_{\theta^{-1}_\ell}(d^\theta_{\theta^{-1}}), \tag{7}$$

$$u_{\theta^{-1}_h}(\mathbf{0}) = u_{\theta^{-1}_h}(d^\theta_{\theta^{-1}}). \tag{8}$$

Proceeding by induction as above, we construct decoys for all $\theta' < \theta$. Now, we define a solution σ as follows. The single-crossing property ensures that each type $\theta' \geqslant \theta$ chooses their corresponding (decoy) contract out of the menu $\{c_\theta, d^\theta_{\theta_1}, \ldots, d^\theta_{\theta_m}, \mathbf{0}\}$ in frame h. At the root of e_θ, θ will choose its contract since it satisfies P^ℓ_θ and any $\theta' > \theta$ will choose the outside option. The $\{d^\theta_{\theta'}\}_{\theta'<\theta}$ are not attractive to those types by single crossing. Turning now to the types $\theta' < \theta$, single crossing also ensures that they prefer the outside option or c_θ at the terminal stage over the decoys. At the root of e_θ, they choose their respective decoy by construction and single crossing. Furthermore, it is less attractive than the outside option in the high frame. We formally verify the construction in Appendix A. □

As is apparent from the above construction, the decoys in the final stage (associated with an increase in payoff-type) circumvent the downward IC constraints and the decoys in the intermediate stage (associated with a decrease in payoff-type) circumvent the upward IC constraints. Therefore, a two-stage decision problem with frames h-ℓ would be sufficient if only upward IC are of concern. All three stages are needed even if only downward IC are of concern, as starting in the high frame allows for revealed types with a relaxed participation constraint.[16]

3.4. *Discussion*

Weakening comonotonicity. Our assumptions can be relaxed at the cost of parsimony. Suppose that (1) *there exists a unique highest frame*, i.e. there exists $h \in F$ such that $\{h\} = \bigcap_{\theta\in\Theta} \text{argmax}_{f\in F}\theta_f$; and (2) *comonotonicity holds locally*, i.e. for each θ there exists a "second highest" frame $\ell(\theta) \in \text{argmax}_{f\in F\setminus\{h\}}\theta_f$ such that for all lower types $\theta'_h < \theta_h$ we have $\theta'_{\ell(\theta)} \leqslant \theta_{\ell(\theta)}$ and $\theta'_{\ell(\theta)} < \theta'_h$ and for all higher types $\theta_h < \theta'_h$ we have $\theta'_{\ell(\theta)} < \theta'_h$. Then our results generalize, replacing ℓ with $\ell(\theta)$ when constructing the continuation problem e_θ.[17] A simple sufficient condition for (1) and (2) is that there is an unambiguously highest and second highest frame.

Consider the following example that violates Assumption 1 but satisfies the assumption above. A product has n flaws and there are n types of consumers, such that for type θ^i flaw i is irrelevant. The sales person can either avoid discussing the flaws (high frame), or focus the attention on one of them (frame l_i). The principal can implement any contract satisfying the participation constraints in the highest frame using a canonical EDP with $\Theta_C = \Theta$ and the type-specific low frame in the second stage. In the screening application of the following section, the principal can extract all surplus in this case.

Commitment and direct mechanisms. The agent is sophisticated but lacks commitment. This is crucial, as the power of the principal to relax IC constraints by concealing types relies on the resulting dynamic inconsistency. In particular, this implies that our contracts cannot be implemented by a direct mechanism. Restricting to direct mechanisms effectively gives commitment as a single-stage interaction does not allow for dynamic inconsistencies. As observed by Galperti (2015), with dynamically inconsistent agents the revelation principle does not apply directly. Instead, agents need to resubmit their complete private information at every stage. In our setting, working with indirect mechanisms is more convenient.

Random mechanisms. We restrict the planner to a deterministic extensive-form mechanism. This is without loss for the implementation of a deterministic contract. To see why, note that the

16. This relates to the findings of Esteban and Miyagawa (2006a), who study a screening problem with temptation preferences when the principal designs a two-stage EDP (menu of menus). As in Section 4, it is the downward constraints which are essential to circumvent. The principal achieves this and fully extracts the untempted surplus when consumers are tempted to overconsume (corresponding to ℓ–h) but cannot circumvent the IC constraints when consumers are tempted to underconsume (corresponding to ℓ–h). Similarly, Yu (2020) shows that the full surplus of the patient period 0 self can be extracted in a β–δ setting with immediate consumption and delayed payment (period 0 and period 1 preferences corresponding to ℓ–h).

17. For a given contract, it is sufficient that local comonotonicity holds for all concealed types. Furthermore, if $\theta'_{\ell(\theta)}$ is sufficiently low for $\theta' > \theta$—violating Assumption 1—it might be that an explicit decoy is not even required.

only way a random mechanism could weaken the conditions for implementation is by relaxing the P^ℓ_θ constraints for concealed types. If the contract violates this constraint for a concealed type, however, this type cannot pass through a decision node with frame ℓ with any positive probability. But then the type is revealed and randomization cannot help implementation.

Due to the discrete nature of moving from revealing to concealing a type, it can be advantageous to implement a random contract using a random mechanism even for a principal with a strictly concave objective. We return to this point at the end of the next section.

Participation constraint at every stage. We assume that the agent can opt-out and choose the outside option at every stage of the decision problem. This is crucial for the trade-off between relaxing participation (by revealing) and relaxing incentive compatibility (by concealing). A weaker restriction would be to require the outside option to be available only at the root. This corresponds to a mechanism that agents enter voluntarily already knowing their type but cannot exit at will. In this case, the principal can implement any vector of contracts that satisfies the participation constraint in the highest frame. The structure (three stages and two frames) and construction of decoys are analogous, but concealing a type no longer involves a tightening of the participation constraint.

Anticipated and lingering effects of frames. We assume that the principal is unrestricted in the choice of frames and, in particular, that a change of framing is effective.[18] One might suppose that framing effects are instead partially "sticky". That is, if the principal is choosing f' after f, then the agent's payoff type will be $\alpha\theta_{f'}+(1-\alpha)\theta_f$ for some $\alpha \in (0.5, 1]$. Similarly, suppose that the effect of the frames used through the decision problem affect the agents evaluation at every stage, as they are considered by the agent when parsing the problem as a whole and stay in the back of the mind of the agent. That is, the agent's payoff type when put in frame f is given by $\alpha\theta_f+(1-\alpha)\tilde{\theta}$ where $\tilde{\theta}$ is the average payoff type over all frames used in the decision problem. Our results generalize to both cases.

Contracts below the outside option. The main theorem generalizes naturally to contracts that may contain negative outcomes (i.e. outcomes smaller than the outside option). When all are in the negative domain, single-crossing implies that lower payoff types have higher valuations. Accordingly, the characterization is a mirror image of Theorem 1 simply replacing the highest and second-highest with the lowest and second-lowest frame. When the vector of contracts has both negative and positive outcomes, the frame at the root becomes a free variable. Such a vector is implementable if and only if it is implementable in a canonical EDP with *some* frame at the root. The structure stays the same, but there is one novel feature: If the contract satisfies the participation constraint for all types in an intermediate frame (i.e. one that is neither the highest nor the lowest), there is no trade-off between incentive compatibility and participation. Using a frame higher than the frame at the root for types with a positive outcome and a frame lower than the frame at the root for types with a negative outcome, the principal can conceal all types without tightening their participation constraint.

Binding bounds on quantities. The canonical EDP gains its implementation power from its decoys, which can be constructed for any vector of contracts and any comonotonic vector

18. This is in line with evidence showing that framing effects, such as gain-loss, are observed within-subject (Tversky and Kahneman, 1981) and even among philosophers who claim to be familiar with the notion of framing, to have a stable opinion about the answer to the manipulated question and were encouraged to consider a different framing from the one presented (Schwitzgebel and Cushman, 2015).

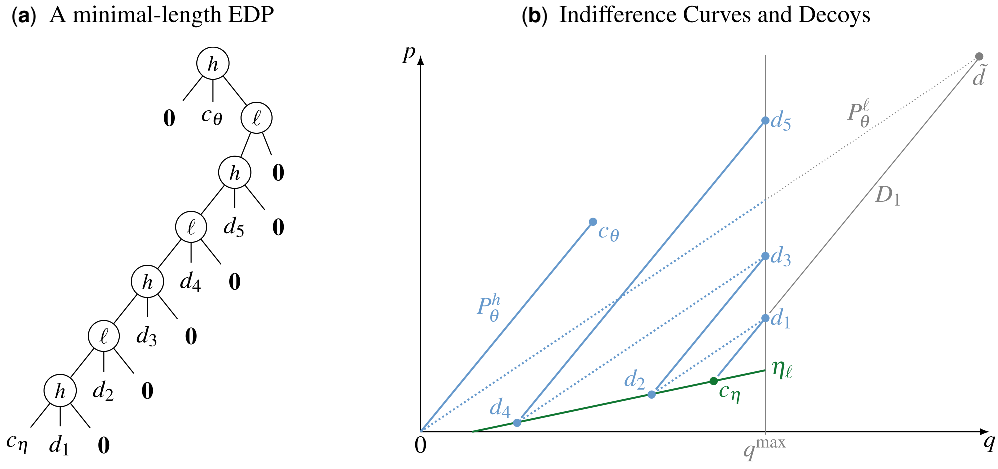

FIGURE 4
Example 2.

of types and frames. The construction, however, does require sufficient room in the space of outcomes to accommodate these decoys. Figuratively speaking, constructing decoys is akin to parallel parking: You can park a car no matter its size and steering angle (corresponding here to the contracts and payoff types) in one swoop, as long as the parking space is sufficiently long. For any fixed length, however, it may require many back-and-forth manoeuvres, the exact number of which depends on these details. Similarly, with a restricted space of contracts there can be a vector of contracts that is implementable by a k-EDP but not by a shorter EDP, for arbitrary odd k, as the following example illustrates.

Example 2 *Consider a setting with two types, $\theta > \eta$, two frames and linear valuation $v_{\theta_f}(q) = \theta_f q$. Now suppose $\mathcal{Q} = [0, q^{\max}]$, where we take the upper bound to be binding in constructing the decoy. Consider a contract $c_\theta \succcurlyeq_{\theta_h} \mathbf{0}$ and $c_\eta \succcurlyeq_{\eta_\ell} \mathbf{0}$ and assume that the upward IC is slack.*

Analogous to Lemma 1, we need to construct decoys (Figure 4). The decoy at the terminal stage of the continuation problem for η must be north-west of the line η_ℓ and south-east of the line D_1. In order to reduce the deviation surplus of the high type as much as possible, it is best to choose the highest quality decoy, i.e. d_1, and move to d_2 in the intermediate stage (frame ℓ): This is the contract with the lowest surplus from the point of view of θ_h that is unattractive to η_ℓ and acceptable over the previously chosen decoy for θ_ℓ. Clearly, this two-stage continuation problem is not sufficient to reduce the deviation surplus to zero. This is because we could not push d_1 to the intersection $\tilde{d}$ of the line D_1 and P^ℓ_θ due to the upper bound on q. Having reduced the deviation surplus of the high type as much as possible at this stage, we can however proceed iteratively until we can place a decoy north-west of the line P^ℓ_θ. Then, the high type opts out in frame ℓ and we have thus reduced his deviation surplus to zero and implemented the contracts.

We formalize this construction in the appendix and show that at least $2\lceil m \rceil$ steps are required to implement the contract, where $m = \log\left(\frac{q^{max}}{q^{max}-q_0} + \frac{q_0\eta_\ell - p_0}{(q^{max}-q_0)(\theta_\ell - \eta_\ell)}\right) / \log\left(1 + \frac{\theta_h - \theta_\ell}{\theta_\ell - \eta_\ell}\right)$ and $(p_0, q_0) := c_\eta$. As is easy to see, m grows without bounds as $q^{max} \searrow q_0$, as c_η delivers more rent and as the impact of framing on the high type vanishes. It shrinks to 1 as q^{max} grows and as the effect of framing on the high type increases.

4. APPLICATION: OPTIMAL SCREENING

This section applies the insights developed in the characterization of the implementable contracts to the monopolistic screening problem. We build on the classic model of price discrimination (Mussa and Rosen, 1978; Maskin and Riley, 1984). As in the implementation problem, the firm designs an extensive-form decision problem and for every decision node picks a frame affecting consumer valuations.

4.1. *The firm's problem*

The monopolist produces goods of quality in $\mathcal{Q}=[0,\infty)$ at a convex cost $\kappa : \mathcal{Q}\to\mathbb{R}$ which is twice-differentiable and satisfies regularity conditions: $\kappa(0)=0$, $\kappa'>0$, $\kappa''>0$ and $v'_{\theta_f}(0)-\kappa'(0)>0$, $\lim_{q\to\infty} v'_{\theta_f}(q)-\kappa'(q)<0$ for all $\theta_f\in\mathbb{R}_{++}$. We denote the efficient quality for a payoff type θ_f by $\widehat{q}_{\theta_f}$ which is defined as the solution of

$$v'_{\theta_f}(\widehat{q}_{\theta_f})=\kappa'(\widehat{q}_{\theta_f}). \tag{9}$$

The efficient quality is unique, positive and strictly increasing in θ_f by our assumptions on v and κ. We denote the contract offering this quality and extracting all surplus from the corresponding payoff type by $\widehat{c}_{\theta_f} := \left(v_{\theta_f}(\widehat{q}_{\theta_f}), \widehat{q}_{\theta_f}\right)$.

A remark on the assumption $\mathcal{Q}=[0,\infty)$ is in order. It corresponds to a situation in which goods produced by the firm are superior in quality to the outside option (normalized to 0) which is interpreted as not purchasing at all. While this is a natural assumption in this context, it also precludes the direct application of Theorem 1 to show that all contracts—and hence *a fortiori* the optimal ones—are canonically implementable.

Given a vector of contracts $\mathbf{c}=\left((p_\theta,q_\theta)\right)_{\theta\in\Theta}$, the profit of the firm is given by[19]

$$\Pi(\mathbf{c}) := \sum_{\theta\in\Theta} \mu_\theta\,(p_\theta-\kappa(q_\theta)),$$

where $\mu_\theta\in(0,1)$ is the prior probability of type $\theta\in\Theta$. Finally, the firm designs an EDP to maximize profits

$$\Pi^* := \sup_{e\in\mathcal{E},\mathbf{c}\in\Sigma^e} \Pi(\mathbf{c}). \tag{Opt}$$

We say that $\mathbf{c}$ (canonically) solves (Opt) if it is (canonically) implementable and $\Pi(\mathbf{c})=\Pi^*$.

The firm's problem is animated by the interaction of framing and extensive forms. Only both features together allow the principal to use different frames at different stages of the decision and thereby generate violations of dynamic consistency which can be exploited. If the firm were restricted to choose a 1-EDP, i.e. a menu and a frame, or if only a single frame were available, the problem collapses to a simple static screening problem (Observation 1). In this case, it is optimal to choose the highest frame h, in order to maximize consumer valuations.

19. In particular, we assume that there are no direct costs of the sales interaction. Our results are qualitatively robust to this possibility. If there are costs per frame, the principal still uses at most two frames, but which ones will depend on their relative costs. If the number of stages itself is the source of the costs, the solution may be a two-stage EDP.

4.2. *Optimal contracts*

We now show that the principal's (Opt) problem over the space of extensive forms is equivalent to a two-step maximization problem based on the necessary conditions for implementation. This relaxed problem characterizes the optimal vector of contracts.

An equivalent problem in price-quality space Consider the profit-maximization problem over contracts subject to the necessary condition for implementation (Proposition 1). Recall that these conditions are indexed by the set of concealed types which is now an additional choice variable for the principal. Clearly, this problem is a relaxation of (Opt).

$$\Pi^R = \max_{\Theta_C \subseteq \Theta} \sup_{\mathbf{c} \in \mathcal{C}} \Pi(\mathbf{c}) \tag{RP}$$

$$\text{s.t.} \quad u_{\theta_h}(c_\theta) \geqslant 0, \quad \forall \theta \in \Theta_R := \Theta \setminus \Theta_C \tag{P^h_θ}$$

$$u_{\theta_\ell}(c_\theta) \geqslant 0, \quad \forall \theta \in \Theta_C \tag{P^ℓ_θ}$$

$$u_{\theta_h}(c_\theta) \geqslant u_{\theta_h}(c_{\theta'}), \quad \forall \theta \in \Theta, \theta' \in \Theta_R \tag{$\mathrm{IC}^h_{\theta\theta'}$}$$

We say that $\mathbf{c}$ solves (RP) if (RP) has a solution $(\Theta_C, \mathbf{c})$.

In this application, Theorem 1 is not directly applicable and not *every* vector of contracts satisfying the necessary conditions is implementable, as we cannot guarantee a sufficiently large space of outcomes. Nevertheless, the *optimal* contract subject to the necessary conditions is implementable.

Theorem 2 *For each $\mathbf{c} \in \mathcal{C}$, the following statements are equivalent:*

(i) $\mathbf{c}$ *solves* (Opt)*,*
(ii) $\mathbf{c}$ *canonically solves* (Opt)*,*
(iii) $\mathbf{c}$ *solves* (RP)*.*

Moreover, such a solution exists and the set of concealed types Θ_C in (ii) can be taken to be same as Θ_C in (iii) and vice versa.

In other words, the optimal contract is canonically implementable and characterized by the relaxed problem (RP). The key step of the proof of Theorem 2 is to show that in the solution of (RP), all upward IC into concealed types are satisfied. Therefore, in order to implement this solution in a canonical EDP, we only need to discourage downward deviation by placing decoys in the final stage of all continuation problems. This construction does not require qualities below the outside option and is hence possible in the present setting.

Without the equivalent formulation in (RP), even verifying the existence of a solution to (Opt) can be troublesome. Theorem 2 shows that instead of a complex optimization problem defined over extensive forms, the principal can solve well-behaved contracting problems over a menu of price-quality pairs, one for each potential set of concealed types and compare the attained values to find the optimum.[20] Once the principal found the (RP)-optimal concealed types and vector of contracts, it is easy to construct a canonical EDP implementing it using Lemma 1.

20. Indeed, a stronger result holds. The (RP)-optimal contracts for any, even sub-optimal, set of concealed types is implementable. The problem can be further simplified by noting that only local IC—those into the nearest revealed types—are binding. See Appendix A.

This result implies that the optimal sales interaction takes a simple form, which we interpret as follows: At the beginning, the agent is presented a range of contracts $\{c_\theta\}_{\theta\in\Theta}$ while the salesperson focuses their attention on quality (high frame). Some of those contracts (those intended for revealed types) can be signed immediately, some others (those intended for concealed types) are only available after an additional procedure that gives the agent some time to consider, while sales pressure is reduced (lower frame). This can be an explicit wait period, where the agent is asked to think about the contract and recontact the seller. Alternatively, the change in frame could be achieved by a change in the salesperson or by acquiring a confirmation that this type of offer is even available for the agent. If the agent is still interested after this ordeal, she is presented with additional offers, the decoy contracts. On path, these offers remain unchosen, the agent chooses the contract she initially intended to obtain.

No shut-down In the classic model of screening, it is sometimes optimal for the monopolist to exclude low types by not selling to them. In our model, this is never the case because concealing a type is always strictly better for the monopolist than exclusion.

Proposition 3 *The optimal contract* (p^*_θ, q^*_θ) *for a type* θ *satisfies* $0 < q_\theta \leqslant q^*_\theta \leqslant \widehat{q}_{\theta_h}$, *where* $v_{\theta_h}(q_\theta) - \kappa(q_\theta) := v_{\theta_\ell}(\widehat{q}_{\theta_\ell}) - \kappa(\widehat{q}_{\theta_\ell})$. *In particular, every type buys a strictly positive quality.*

Indeed, concealing a type can be interpreted as a soft form of shut-down. In order to eliminate information rents, the principal reduces the revenue extracted from a type. The key difference is that it can be achieved at a strictly positive quality, while extracting revenue from this type.

Optimal contracts for concealed types For concealed types, we can provide an additional lower bound on quality in the optimal contract. The contract for concealed types is subject to constraints in two frames: a participation constraint in the lower frame ℓ and an IC constraint in the higher frame h. Since concealed types cannot be imitated, there is no reason to distort their quality downward below the efficient quality in the lower frame, $\widehat{q}_{\theta_\ell}$. It can be optimal, however, to increase the quality above this level in order to deliver rent more cost-effectively in order to satisfy their IC constraint.

Proposition 4 *Consider a concealed type* $\theta \in \Theta_C$. *Then, the optimal quality is bounded between the efficient quality in frame* ℓ *and* h: $\widehat{q}_{\theta_\ell} \leqslant q_\theta \leqslant \widehat{q}_{\theta_h}$. *In particular, the optimal contract is*

$$(p_\theta, q_\theta) = \begin{cases} \widehat{c}_{\theta_\ell}, & \text{if } \Delta_\theta \leqslant v_{\theta_h}(\widehat{q}_{\theta_\ell}) - v_{\theta_\ell}(\widehat{q}_{\theta_\ell}), \\ \big(v_{\theta_\ell}(q^*), q^*\big), & \text{if } \Delta_\theta \in \big[v_{\theta_h}(\widehat{q}_{\theta_\ell}) - v_{\theta_\ell}(\widehat{q}_{\theta_\ell}), v_{\theta_h}(\widehat{q}_{\theta_h}) - v_{\theta_\ell}(\widehat{q}_{\theta_h})\big], \\ \big(v_{\theta_h}(\widehat{q}_{\theta_h}) - \Delta_\theta, \widehat{q}_{\theta_h}\big), & \text{if } \Delta_\theta \geqslant v_{\theta_h}(\widehat{q}_{\theta_h}) - v_{\theta_\ell}(\widehat{q}_{\theta_h}), \end{cases} \tag{10}$$

where q^* *solves* $v_{\theta_h}(q^*) - v_{\theta_\ell}(q^*) = \Delta_\theta$, *and* $\Delta_\theta := \operatorname{argmax}_{\theta' \in \Theta_R} u_{\theta_h}(c_{\theta'})$ *denotes the rent delivered to type* $\theta \in \Theta_C$, *and* $\mathbf{c}$ *is the optimal contract.*

If the required rent is low, only the participation constraint in the low frame binds and the optimal contract is the efficient contract for the type in the low frame. As more rent needs to be delivered in the high frame, it becomes optimal to increase the quality of the product up to the efficient quality in the high frame.

The contract further illustrates the cost of concealing a type. From the perspective of the high frame, a concealed type receives at least the minimal rent $v_{\theta_h}(\widehat{q}_{\theta_\ell}) - v_{\theta_\ell}(\widehat{q}_{\theta_\ell})$, reducing the payoff

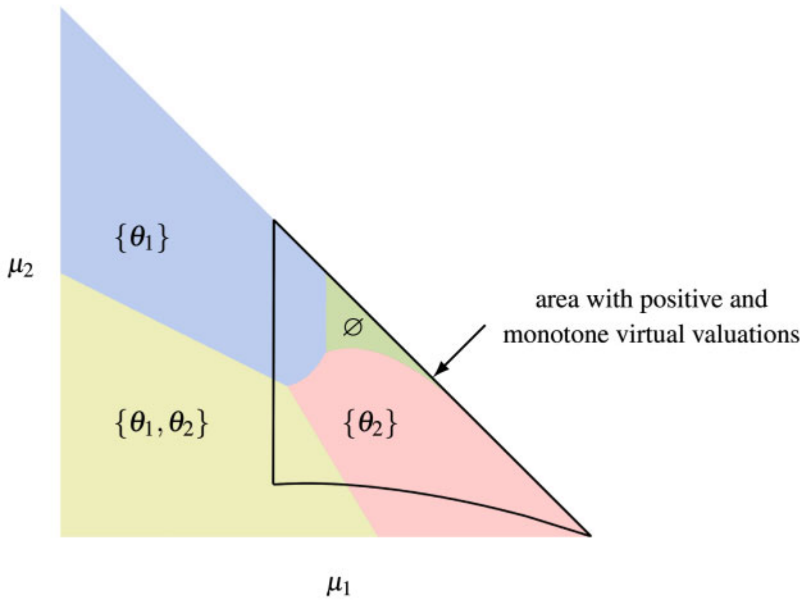

FIGURE 5
Optimal Θ_C for $\theta^1=(1,3),\theta^2=(4,5),\theta^3=(5,6)$.

of the principal. The cost of concealing a type is decreasing in the information rent Δ. If Δ is sufficiently high (in the third regime of (10)), it is costless to conceal the type.

4.3. *Optimal concealed types*

Analogous to the distortion and exclusion in the classic screening problem, one might conjecture that it is optimal for the principal to conceal low types and reveal high types. This is not true in general but contains a grain of truth: Types are concealed in order to eliminate downward deviations into them, which is not a concern for the highest type. Revealing the highest type is therefore always optimal.

Observation 2 *Suppose* $(\Theta_C^*,\mathbf{c}^*)$ *solves* (RP) *and the highest type* $\overline{\theta}=\max\Theta$ *is concealed, i.e.* $\overline{\theta}\in\Theta_C^*$. *Then* $(\Theta_C^*\setminus\{\overline{\theta}\},\mathbf{c}^*)$ *also solves* (RP).

In general, there are no other restrictions on the optimal set of concealed types, as the following linear-quadratic three-type example illustrates. In Figure 5, we plot the regions of the probability simplex where particular sets of concealed types are optimal. All four remaining cases are realized for some distribution. In addition, the restriction to monotone virtual values, which ensures monotonicity in the classic screening model, does not rule out any configuration.

Loosely speaking, concealed types are substitutes for the principal. Consider two types $\theta<\theta'$. By concealing θ, the principal reduces the rent θ' obtains, increasing the costs of concealing θ' (as it is more costly to conceal a type if it has a low information rent; Lemma 4). In addition, a lower rent implies that concealing θ' has a smaller benefit as well, as information rents compound. Similarly, concealing θ' reduces the benefit of concealing θ. This pattern of substitutability is reflected in Figure 5 as the regions $\Theta_C=\{\theta_1\}$ and $\Theta_C=\{\theta_2\}$ touch.

Sufficiently likely types are revealed. It is not profitable to conceal very likely types, since the gain from the reduction of information rents for other types is outweighed by the loss of profits that can be extracted from them directly.

Proposition 5 *For any type θ there exists a probability threshold $\bar{\mu}_\theta \in (0,1)$, such that for any $\mu_\theta \in [\bar{\mu}_\theta, 1]$, an optimal set of revealed types contains θ.*

This proposition suggests interpreting the contracts of revealed types as standard options, which are relevant for common types of consumers and available immediately in the store, and the contracts for concealed types as specialty options relevant for rare types of consumers and available only on order.

High θ_ℓ favours concealing. The difference between the valuations in frames h and ℓ determines the cost of concealing. If we fix all payoff types but *increase* the ℓ-frame valuation of a concealed type, this cost is reduced and this type remains concealed.

Proposition 6 *Let Θ_C be an optimal sets of concealed types for (Θ, μ) and let $\theta \in \Theta_C$. Define $\tilde{\theta}$ such that $\tilde{\theta}_\ell \geqslant \theta_\ell, \tilde{\theta}_h = \theta_h$. Then, for the set of types $(\Theta \setminus \{\theta\}) \cup \{\tilde{\theta}\}$ there exists a solution of the principal's problem (RP) with the set of concealed types $\tilde{\Theta}_C := (\Theta_C \setminus \{\theta\}) \cup \{\tilde{\theta}\}$.*

Fixing the highest valuation the principal can achieve for each type, the cost of concealing is low if θ_ℓ is high. We can interpret this as a more precise control of the principal over consumer valuations. With sufficient control, she will conceal all types except for the highest.

Proposition 7 *For any Θ, there exists an $\varepsilon > 0$ such that for any type space $\tilde{\Theta}$ with $\{\theta_h\}_{\theta \in \Theta} = \{\theta_h\}_{\theta \in \tilde{\Theta}}$ and $\max_{\theta \in \tilde{\Theta}}(\theta_h - \theta_\ell) < \varepsilon$, the optimal set of concealed types is $\tilde{\Theta} \setminus \{\max \tilde{\Theta}\}$.*

4.4. *Extension to naive consumers*

So far, we have only considered fully sophisticated agents. We now study naive consumers. They understand the structure of the extensive-form decision problem and the choices available to them, but they fail to anticipate the effect of framing. Faced with an EDP, they pick the continuation problem containing the contract they prefer most in their current frame.[21] They fail to take account of the fact that in this continuation problem, they may be in a different frame and end up choosing a different contract.

Setup Towards the definition of a naive solution, let $\mathcal{C}(e)$ denote the set of contracts in an EDP e. That is, letting $\mathcal{C}(e^0) = e^0$ for $e^0 \in \mathcal{E}^0$, define

$$\mathcal{C}(e) := \bigcup_{e' \in E} \mathcal{C}(e'), \text{ for } e = (E, f).$$

21. A related idea is projection bias (Loewenstein, O'Donoghue and Rabin, 2003). The main difference is that our construction depends on the consumers' ability to forecast their future actions, not tastes. In this general sense, sophisticated consumers exhibit no projection bias, while naive consumers exhibit complete projection bias.

Now call $s_\theta: \mathcal{E}\cup\mathcal{E}^0 \to \mathcal{E}\cup\mathcal{E}^0$ a naive strategy for θ if for all $e^0\in\mathcal{E}^0$, $s_\theta(e^0)=e^0$, for all $(E,f)\in\mathcal{E}$, $s_\theta(E,f)\in E$, and

$$\mathcal{C}(s_\theta(E,f))\cap \underset{\mathcal{C}(e)}{\operatorname{argmax}} u_{\theta_f} \neq \varnothing.$$

Put differently, when facing $e=(E,f)$, a consumer identifies the f-optima in the set of all contracts in e, $\mathcal{C}(e)$, and chooses a continuation problem containing an optimum.

We call $\nu: \Theta \to \mathcal{C}$ a naive solution of an EDP e if there exists a naive strategy profile s such that any type θ arrives at $\nu(\theta)$ by following s_θ, i.e. $\nu(\theta)=(\underbrace{s_\theta \circ \cdots \circ s_\theta}_{k \text{ times}})(e)$ for $e\in\mathcal{E}^k$. Let N^e be the set of all naive solutions to an EDP e.

We consider the case when there are both naive and sophisticated consumers and the principal cannot observe their cognitive type. Let $\Theta=\Theta_S\sqcup\Theta_N$ be the disjoint union of the set of sophisticated types Θ_S and the set of naive types Θ_N. That is, we allow for the existence of $\theta^s\in\Theta_S$ and $\theta^n\in\Theta_N$ that differ only in their sophistication, but not in their tastes conditional on any frame. Define the optimal profits similarly to (Opt) as

$$\begin{aligned}\Pi^* := \max_{e\in\mathcal{E},\mathbf{c}\in\mathcal{C}} &\Pi(\mathbf{c}) \\ \text{s.t. } &c_\theta \in \Sigma^e(\theta), \forall\theta\in\Theta_S, \\ &c_\theta \in N^e(\theta), \forall\theta\in\Theta_N.\end{aligned} \tag{11}$$

Optimal structure and contracts We illustrate in an example how the principal can use decoys to screen when naive types are present.

Example 3 *Recall from Example 1 that there are two frames, $\{\ell,h\}$, and two payoff types, $\{\theta^1,\theta^2\}$. The key construction can be illustrated using three equally likely types, two naive and one sophisticated. There is a naive version of both payoff types, and a sophisticated high type, formally $\Theta=\Theta_S\sqcup\Theta_N=\{\theta^{s2}\}\sqcup\{\theta^{n1},\theta^{n2}\}$. In this setting, the principal can sell the h-efficient quality to naive consumers and fully extract their surplus. This creates no information rents for the sophisticated type—screening by cognitive type is free. As a result, she can also implement the high-frame full-extraction contract for θ^{s2}. The optimal EDP is given in Figure 6. It implements $c^{n1}=(16,4)$, $c^{n2}=(36,6)$, $c^{s2}=(36,6)$.*

First, consider the sophisticated type. As in Example 1, the contracts c^{n1},b^{n2} are more attractive than the implemented c^{s2}, but are concealed using the decoys $d^{s2},d^{s2\prime}$, respectively.[22]

Let's turn to the naive types. The leftmost continuation problem is intended for θ^{n1}. Even though θ^{n1} is concealed, the principal extracts full surplus in the high frame. How is this possible? At the second stage in frame ℓ, he indeed prefers the outside option over c^{n1}. But, he wrongly believes that he will choose the outside option after continuing. Hence, he continues and—back in frame h—chooses c^{n1}.

In order to implement the contract for θ^{n2}, the principal needs to use a decoy. At the root, he strictly prefers c^{n1} to c^{n2}. In order to lure him into the middle continuation problem, the principal introduces a decoy b^{n2}. This decoy works differently from the decoys used with sophisticated consumers. It serves as bait and is the most preferred contract out of the whole decision problem

22. In this simple example, the three decoys can coincide, $d^{s2}=b^{n2}=d^{s2\prime}=(40,8)$, because there are only two different payoff types.

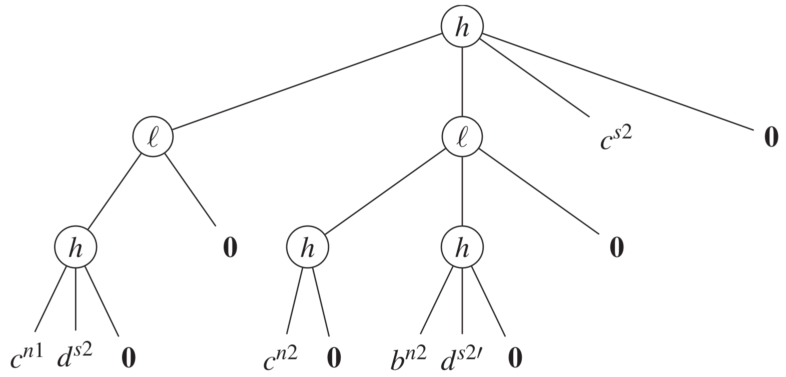

FIGURE 6

The optimal extensive-form decision problem in Example 3.

for θ^{n2}. As a consequence, he continues into the middle continuation problem. There, the switch happens: b^{n2} is unattractive from the perspective of the low frame and θ^{n2} continues, expecting to pick the outside option in the continuation problem. Like θ^{n1} he reconsiders at the terminal node and ends up with c^{n2}.

This construction generalizes.[23] The optimal EDP achieves the same solution as if the principal knows which consumers are naive and the types of the naive consumers. Naive types do not receive information rents, they obtain the full extraction contract in the high frame. Sophisticated consumers obtain the optimal contract according to Theorem 2.

Theorem 3 *Let σ and ν be firm-preferred sophisticated and naive solutions of an optimal EDP, respectively. Then $\nu(\theta)=\widehat{c}_{\theta_h}$ for all $\theta \in \Theta_N$ and σ is a firm-preferred optimal sophisticated solution for the set of types Θ_S and the conditional prior. Moreover, there exists an optimal EDP with the high–low–high structure.*

The principal also uses decoy contracts for naive consumers, but their role is reversed: In the construction for sophisticated consumers, we placed decoys in continuation problems to make sure that no other type wants to enter the continuation problem, as they correctly anticipate that they would choose the decoy. In the case of naive consumers, instead of decoys to repel imitators, we introduce decoys in order to lure types into their corresponding continuation problems. Agents wrongly believe that they will choose their respective decoy, which is the most attractive contract in the whole EDP for them in their current frame. Once types are separated at the root of the decision problem, the dynamic inconsistency introduced by changing frames allows the decision problem to reroute consumers from their decoy to the intended contract.

The optimal extensive-form decision problem retains the simple three-stage structure, we only add a continuation problem for each naive type to the extensive form described in Theorem 1.

23. Without principal-preferred tie breaking, the first-best in the highest frame can be implemented virtually. For the example, the principal places an additional bait decoy together with c^{n1} (c^{n2}, resp.). This bait gives the respective type a rent of ε from the perspective of frame ℓ. Therefore, they strictly prefer to continue at the intermediate stage. To obtain strict incentive compatibility in the final stage—both relative to the outside option and to the newly added bait decoys—the contracts c^{n1} and c^{n2} (and the decoys) need to be perturbed resulting in rents of order ε. As $\varepsilon \to 0$, the solution converges to the solution with principal-preferred tie breaking. This virtual implementation result holds in general for Theorem 3.

Consequently, the optimum can be achieved by a three-stage EDP with $|\Theta|$ continuation choices at the root, similar to a canonical EDP. As in the optimal EDP for sophisticated consumers, we do not require second-period decoys. In the continuation problem for naive consumers, we have a decoy menu that lures in naifs while ensuring that sophisticated types cannot deviate to the decoy. We provide the details of the construction in Appendix A.

Welfare gains from sophistication Are consumers better off if they are sophisticated? Welfare statements in the presence of framing are generally fraught with difficulty. In this case, we can rank the contracts obtained by sophisticated and naive agents from a consumer perspective without taking a stand on the welfare-relevant frame. In the following sense, sophistication partially protects consumers from exploitation through the use of framing.

Observation 3 *For all types, the contract under sophistication is weakly preferred to the contract under naivete from the perspective of every frame.*[24]

From an efficiency perspective, the two cases are not unambiguously ranked. For naive consumers, the principal implements the efficient quality from the perspective of the highest frame. Quality is lower for sophisticated consumers, an efficiency gain from the perspective of all frames except the highest one.

Discussion: partial naivete We can also extend our results to partial (magnitude) naivete. Denote the parameter determining the intensity of naivete by $\alpha \in [0,1]$, with $\alpha = 0$ representing full sophistication. Suppose a consumer with current payoff type θ_f anticipates a future choice that will actually be made according to payoff type $\theta_{f'}$. Let $\widehat{\theta}(\theta_f, \theta_{f'}, \alpha)$ denote what he currently perceives to be his future payoff type. Assume $\widehat{\theta}$ is increasing in the first two arguments, monotonic in α and satisfies $\widehat{\theta}(\theta_f, \theta_f, \alpha) = \theta_f$ for all α. Under full sophistication we have $\widehat{\theta}(\theta, \theta', 0) = \theta'$, under full naivete $\widehat{\theta}(\theta, \theta', 1) = \theta$. This structure ensures that a partially naive agents' predictions satisfy comonotonicity (Assumption 1).

Both the sophisticated construction (using decoys as poison pills) and the naive construction (using decoys as bait) generalize to partial naivete: An arbitrarily small degree of sophistication (respectively naivete) is sufficient. Since the naive construction achieves full extraction, it is preferable for the principal. Therefore, there is a discontinuity in the optimal mechanism and profit at full sophistication.[25] Formally, this occurs because the quality of the decoys used in the mechanism grows without bounds as $\alpha \to 0$. As a matter of practicality, we therefore would not expect the naive construction to be used for almost fully sophisticated consumers. Instead, the sophisticated construction or an intermediate approach using decoys both as bait and poison pills to extract more surplus could be employed.

4.5. *Additional participation constraints and cool-off regulation*

In many jurisdictions, regulation mandates a right to return a product for an extended period of time after the purchase. The express purpose of such regulation is to allow consumers to cool off

24. This implies a weak improvement in the sense of Bernheim and Rangel (2009) if the two contracts are not identical.

25. A similar observation holds in Gottlieb and Zhang (2021), where the contract for any level of partial naivete converges to the efficient one but the sophisticated contract stays inefficient.

(a) An EDP e **(b)** The interim modification $\bar{e}$ **(c)** The ex-post modification $\underline{e}$

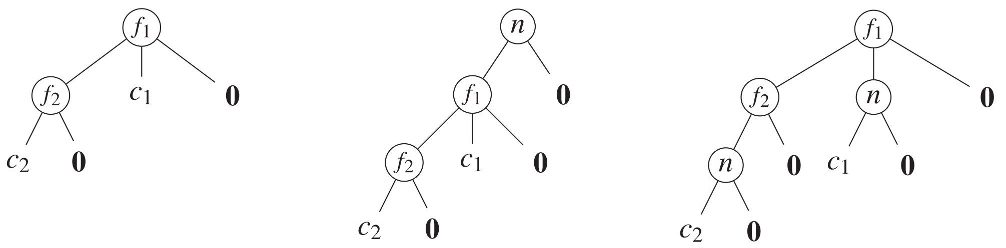

FIGURE 7
Interim and ex post participation constraints in frame n.

and reconsider the purchase in a calm state of mind unaffected by manipulation by the seller.[26] Interestingly, such legislation typically applies only to door-to-door sales and similar situations of high sales pressure to which consumers did not decide to expose themselves. If consumers decide to enter a store or contact a seller, they are not protected by the law. This suggests that legislators consider the option to avoid the firm's sales pressure entirely to protect consumers sufficiently.[27] We evaluate this intuition using our model.

Consider a situation when consumers decide in a neutral frame whether to go to the store. One can interpret this decision as an additional interim participation decision at the root. Alternatively, suppose that there is a regulation that allows consumers to return a product if they wish to do so *ex post* in the neutral frame (as in Salant and Siegel, 2018). One can interpret this decision as an additional ex post participation decision at every terminal decision stage.

Formally, denote the neutral frame by $n \in F$, $n < h$. This is the frame the consumer is in when unaffected by direct sales pressure by the firm.[28] We call $\bar{e} := (\{e, \mathbf{0}\}, n)$ an *interim modification*[29] of e. Next, we define an *ex post modification* $\underline{e}$ of an EDP e by replacing every non-$\mathbf{0}$ contract c with a a 1-EDP $(\{c, \mathbf{0}\}, n)$ recursively. First, for any $e \in \mathcal{E}^0$, let $\underline{e} := \bar{e}$. Then, having defined an *ex post* modification on $\mathcal{E}^j, \forall j = 0, \ldots, k$, we define the ex post modification for any $e = (E, f) \in \mathcal{E}^{k+1}$ as $\underline{e} := (\{\underline{e}'\}_{e' \in E}, f)$. We say that e is an *EDP with an interim (ex post) participation constraint* if it is an interim (ex post) modification of some EDP.

Sophisticated consumers. If consumers are sophisticated, both constraints are equivalent and imply that if a contract is chosen by type θ, then it must satisfy the additional participation constraint P^n_θ. Therefore, the firm implements the efficient allocation associated with frame n and leaves no information rent to consumers.

Observation 4 *Suppose $\Theta = \Theta_S$. Let $\bar{e}^*$ and $\underline{e}^*$ be optimal EDPs with interim and ex post participation constraints. Then, their firm-preferred solutions $\bar{\sigma}$ and $\underline{\sigma}$, respectively, are such that for all $\theta \in \Theta$,*

$$\bar{\sigma}(\theta) = \underline{\sigma}(\theta) = \widehat{c}_{\theta_n}.$$

26. E.g. directive 2011/83/EU: "the consumer should have the right of withdrawal because of the potential surprise element and/or psychological pressure".

27. Rights to return are also motivated by giving consumers an opportunity to physically inspect a good they ordered online, while a cooling-off period before ordering the product can help to protect consumers against projection bias. See Michel and Stenzel (2021) for a comparison of these two policy instruments in this context.

28. One possible effect of marketing is influencing this neutral frame, but we do not consider this margin.

29. Here, the notion of interim modification is defined on $\mathcal{E} \cup \mathcal{E}^0 \setminus \{\mathbf{0}\}$. For simplicity, let $\bar{\mathbf{0}} := \mathbf{0}$

This observation follows immediately from Theorem 2. The principal can remove all incoming IC constraints at the cost of an additional participation constraint in a lower frame. As such a constraint is introduced anyway with interim or ex post participation constraints, the principal can conceal all types at no additional cost.[30] Both restrictions protect against overpurchases relative to the preferences in the neutral frame, but neither protects against the extraction of all information rents by exploiting induced violations of dynamic consistency. Overall, whether sophisticated consumers benefit from these restrictions is ambiguous.

In line with the intuition suggested by policy, sophisticated consumers do not require the additional protection of a right to return if they can avoid the interaction with the firm altogether. They correctly anticipate their future actions and hence—given a choice—only interact with a seller, if the result will be acceptable to them from their current frame of reference rendering a right to return in this frame superfluous.

Naive consumers. With naive consumers, we now need to distinguish between an interim choice to initiate the interaction and an *ex post* right to return in the same neutral frame. While a right to return is still effective, naive consumers cannot protect themselves by avoiding the seller.

Observation 5 *Suppose $\Theta = \Theta_N$. Let $\bar{e}^*$ and $\underline{e}^*$ be optimal EDPs with interim and ex post participation constraints. Then, their firm-preferred naive solutions $\bar{\nu}$ and $\underline{\nu}$ satisfy for all $\theta \in \Theta$,*

$$\bar{\nu}(\theta) = \widehat{c}_{\theta_h},$$
$$\underline{\nu}(\theta) = \widehat{c}_{\theta_n}.$$

The intuition underlying the design of regulation does not apply for naive consumers. They are overly optimistic about the outcome of their interaction with the seller. Naive consumers can always be lured in with attractive decoys and consequently the option to avoid the seller is not sufficient to protect them from over-purchasing. In the optimal EDP, all consumers regret the purchase from the perspective of the neutral frame. A right to return even for in-store sales would offer them additional protection.

4.6. *Discussion*

Random contracts. The principal can do strictly better by randomizing within the canonical mechanism and thereby implementing a random contract. Randomization allows the principal to smooth out the concealment of types. To see this, consider a situation with three types where it is optimal to conceal only the intermediate type and the P^{ℓ}-constraint is binding in his contract. Then, the IC constraint from the highest to the intermediate type is slack at the root, as the intermediate type is concealed and the highest type obtains a strictly positive rent (from the IC to the lowest type). Consider a modification of the mechanism where the intermediate type is concealed with probability $1-\varepsilon$ and revealed otherwise, obtaining the contract that is optimal ignoring the IC constraint of the highest type. The uncertainty resolves after the agent makes his decision at the root, but before the frame-change to ℓ. In this mechanism, the highest type still strictly prefers not to imitate the intermediate type at the root if ε is sufficiently small. Furthermore, *ex ante* profit is strictly greater as the "revealed" contract for the intermediate type

30. Salant and Siegel (2018) show that the principal may not use framing when such a constraint is added to the problem of designing a framed menu. In particular, the principal cannot necessarily extract all rents without the use of an extensive form.

is more profitable than concealing it. Allowing for random contracts, however, the problem loses tractability because local IC constraints are in general no longer sufficient.

Commitment to side payments. In the main section, we assume that the consumer can end the interaction with the principal at every step. Suppose instead that the principal can charge an "entry fee" for a continuation problem. For example, in order to pre-order a certain make of a car, the consumer needs to make a deposit, which is forfeit if the consumer does not follow through. Formally, the principal designs an EDP where every non-terminal node is labelled with a payment in addition to a frame. With sophisticated consumers, the principal can implement the efficient outcome and extract all surplus in the highest frame: All types are concealed and are charged $v_{\theta_h}(\widehat{q}_{\theta_h}) - v_{\theta_\ell}(\widehat{q}_{\theta_h})$ to enter their respective continuation problem. This charge does not interact with the decoy construction and ensures that the concealed type does not opt out in the low frame, as part of the price is already sunk at that stage.

The problem for naive consumers is not well defined. This is because side payments allow to construct a money pump in an EDP with alternating high and low frames as follows. In the high frame, the consumer can either opt out and get an ε-value gift, or continue at small fixed charge. In the low frame, the options are the decoy (attractive in the high, but not the low frame), opt out (without a gift), and to continue. In the naive solution, the consumer always continues: in the high frame—expecting to choose the decoy at the next stage; in the low frame—expecting to opt out with a gift in the next stage. Then, as the length of the EDP grows, the principal can extract unbounded profits.

Principal commitment. Without commitment, the principal cannot take advantage of the power to frame to circumvent incentive compatibility. Consider the case without any commitment for her and fix an EDP and it's solution. Suppose a decision node is reached by this solution and is terminal for all types reaching it. The EDP is consistent with sequential rationality of the principal only if this node is set in the highest frame and engages in optimal static screening given the set of types reaching it. But then, the lowest type obtaining a strictly positive q in this menu (which always exists) obtains no rent, and his contract violates the participation constraint in all lower frames. This is possible only if there are no lower frames en route to this decision node. But then, the EDP is equivalent to a menu set in the high frame. The problem without principal commitment is therefore equivalent to the single-stage problem.

5. CONCLUSION

We analyse the effect of framing in a one-dimensional principal-agent model. The principal can frame decisions in several ways, affecting the agent's valuation as expressed by their choices. Such a setting naturally leads to extensive-form decision problems. The principal uses framing not only to relax individual rationality constraints, but mainly to induce dynamic inconsistency and thereby bypass incentive compatibility constraints, despite strategic sophistication. Our main result is that—with a large enough space of outcomes—any implementable contract can be implemented in a canonical extensive-form decision problem with only three stages and two frames. Only some contracts are available immediately, while others are available after the agent's frame is lowered and raised again for the final choice. At the latter two stages, the principal places decoy contracts, which remain unchosen on-path but are designed to render deviations futile.

We apply our results to the classic monopolist screening problem. The optimal contract is implementable in canonical form, even if the principal cannot offer a lower quantity than the outside option (e.g. the outside option corresponds to not buying at all in this product category). This simple extensive form allows the firm to eliminate information rents at the cost of lower surplus and thereby achieve a payoff that is strictly larger than full surplus extraction at all but the highest frame. Even if consumers are protected by a shop-entry decision or right to return the product in an exogenously given neutral frame, they are not protected against the full extraction of their information rents. We also characterize the outcome with naive consumers. The structure of the optimal extensive form and the contracts of sophisticated agents are robust to the presence of naive types. Naive types can be screened without generating any additional information rents, even if the cognitive type itself is also unobserved.

Our analysis suggests several natural extensions. First, we consider a general but stylized model for our application to outline the power of decision-framing design to extract surplus from consumers and its potentially wide-ranging consequences for consumer protection policy. Applying our results in a more detailed model to questions of behavioural industrial organization and regulation while taking into account the limited nature of commitment, questions of competition, and a more realistic interior degree of sophistication, which may very well be influenced by the usage of extensive form mechanisms, as well as other features of real world sales interactions and mechanisms is a promising direction for future research.

Second, we assumed that choice depends on exogenous factors of the presentation (i.e. the frame) that are chosen by the principal but satisfies the axioms of utility maximization given every frame. If framing affects choice through focusing the attention of consumers on certain attributes, for example, we consider the case where these attributes are emphasized by the salesperson or the information material and assume that the properties of the choice set are not relevant. Alternatively, one might assume that the attention of the consumer is affected by the properties of the choice set (such as in the models of focusing (Kőszegi and Szeidl, 2013) and salience (Bordalo *et al.*, 2013)). With such "endogenous framing", a seller could influence decisions by including an option in the choice set that directs the focus more towards quality. We expect that ideas similar to our construction can be applied to this setting. There is an important caveat, however. While in our setting, the frames are fixed for every decision node independently of the agent's type, context effects depend on the choice set, which is generated by backward induction and is hence type dependent. In effect, the frame can be made type dependent. Consequently, screening with menu dependent preferences is a considerably richer setting and left for future research.

Finally, we focus on the single-agent problem. The idea of using changes in framing to induce dynamic inconsistency which is used to bypass incentive compatibility applies more broadly also to mechanism design problems with multiple agents. To illustrate the possibilities, consider to following simple example.

Example 4 *Consider again the payoffs of Figure 1a, but now there are two agents in a single-object independent private values auction. It is easy to check that the maximal profit from using an auction with a fixed frame is* 4.5. *Consider instead the following scheme: Every agent independently goes through a mechanism analogous to Figure 1b. We interpret arriving at c^1 as reporting type θ^1 and analogously for c^2. The object is allocated to the highest reported type (with uniformly random tie-breaking), with a price of $\theta_h^2 = 6$ for the high type and $\theta_\ell^1 = 3$ for the low type. If only one agent arrives at d^2, he obtains the object with probability 1 and pays* 5.25. *The outside option corresponds to opting out of the auction, if both agents arrive at the decoy, we consider both to be opting out. It is easy to check that truth-telling is an equilibrium of this mechanism. It yields a profit of* 5.25, *beating the best fixed-frame auction.*[31]

31. By revenue equivalence, the highest revenue in a static mechanism is the expected second highest value in the high frame, namely $\frac{1}{4}6+\frac{3}{4}4 = 4.5$. For our mechanism, let us first check that "truthtelling" is an equilibrium: For the low

Generally, it may be beneficial in multi-agent settings to use the information gained in the first stage in order to determine the second stage for other agents, and so on, or even use more complex interdependent schemes. Extending our analysis to a more general Bayesian mechanism-design setting is therefore left for future research.

Acknowledgments. We are grateful for valuable discussions and suggestions to Roland Bénabou, Elliot Lipnowski, Stephen Morris, Pietro Ortoleva, Wolfgang Pesendorfer, Doron Ravid, Leeat Yariv, the editor, and the anonymous referees; and for helpful comments to seminar participants at Princeton University, the Econometric Society Summer School 2018 in Singapore, Young Economists Symposium 2018, ESWM 2018, the NOeG Winter Workshop 2018, CMID 20, EC 2021, CSEF Naples, the CRC TR 224 retreat, and participants of the RSF Summer Institute 2018. Ostrizek gratefully acknowledges financial support from the Philip Bennett *79 Public Policy and Finance Fund and funding by the German Research Foundation (DFG) through CRC TR 224 (Project A01 and B02).

Data Availability Statement

The code underlying this article (Figure 5) is available in Zenodo, at https://dx.doi.org/10.5281/zenodo.6915672.

A. PROOF APPENDIX

A.1. *Preliminaries*

Whenever types are indexed by i, we use notation $u^i_f := u_{\theta^i_f}$. For each θ^i_f, let $\succcurlyeq^i_f$ and (sometimes) $\succcurlyeq_{\theta^i_f}$ denote the corresponding preference relation. For any EDP e, let $\mathcal{C}(e)$ denote the set of all contracts available in e. Formally, for $e=(A,f)\in\mathcal{E}^1$, $\mathcal{C}(e):=A$ and, recursively, for $e=(E,f)\in\mathcal{E}^k$, $\mathcal{C}(e):=\cup_{e'\in E}\mathcal{C}(e')$.

First, note that our increasing marginal differences assumption on v implies that the consumer's preferences exhibit the single-crossing property.

Lemma 2 (Single-crossing property) For any two payoff types $\bar\theta,\underline\theta\in\mathbb{R}$, such that $\bar\theta\geqslant\underline\theta$, and contracts $x,y\in\mathcal{C}$, such that $q^y\geqslant q^x$, we have

$$u_{\underline\theta}(y)\geqslant u_{\underline\theta}(x)\Longrightarrow u_{\bar\theta}(y)\geqslant u_{\bar\theta}(x)$$
$$u_{\bar\theta}(y)\leqslant u_{\bar\theta}(x)\Longrightarrow u_{\underline\theta}(y)\leqslant u_{\underline\theta}(x).$$

Proof. Take any $x,y\in\mathcal{C}$, such that $q^y\geqslant q^x$ and $u_{\underline\theta}(y)\geqslant u_{\underline\theta}(x)$. Note that the increasing differences property ($\frac{\partial^2 v}{\partial\theta_f\partial q}>0$) implies that

$$\begin{aligned}
u_{\bar\theta}(y)-u_{\bar\theta}(x)&=v(\bar\theta,y)-v(\bar\theta,x)+p^y-p^x\\
&=\int_x^y\frac{\partial v_{\bar\theta}}{\partial q}(q)\mathrm{d}q+p^y-p^x=\int_x^y\left(\frac{\partial v_{\underline\theta}}{\partial q}(q)+\int_{\underline\theta}^{\bar\theta}\frac{\partial^2 v_{\theta_f}}{\partial\theta_f\partial q}(q)\mathrm{d}\theta_f\right)\mathrm{d}q+p^y-p^x\\
&\geqslant v_{\underline\theta}(y)-v_{\underline\theta}(x)+p^y-p^x=u_{\underline\theta}(y)-u_{\underline\theta}(x)\geqslant 0.
\end{aligned}$$

The proof of the second implication is analogous. □

Second, we prove the following result which ensures the existence of suitable decoy contracts.

Lemma 3 For any two payoff types $\underline\theta<\overline\theta$, the function $\phi_{\overline\theta,\underline\theta}:=v_{\overline\theta}-v_{\underline\theta}:\mathbb{R}\to\mathbb{R}$ is twice differentiable, strictly increasing, and bijective.

type, the only play that yields a positive expected utility is continuing to c^1 (i.e. reporting his true type). For the high type, he would choose the decoy contract at the final stage after a deviation, as obtaining the object for sure (assuming equilibrium truthful play by the other agent) at a price of 5.25 gives surplus $0.75\geqslant\frac{1}{2}\frac{1}{2}3$, where the latter is the expected surplus from imitating the low type and winning the tie-break. Anticipating this in the intermediate stage, the high type chooses to opt out as the price he expects to pay for the object is higher than his value: $5.25>\theta^2_\ell=5$. Therefore, at the root, no effective deviation is possible and the high type is willing to choose c_2, i.e. to report the type truthfully. The revenue is therefore $\frac{3}{4}6+\frac{1}{4}3=5.25$ as the mechanism extracts all high-frame surplus from the high type and all low-frame surplus from the low type.

Proof. First, $\phi_{\overline{\theta},\underline{\theta}}$ is twice differentiable since so are $v_{\overline{\theta}}$ and $v_{\underline{\theta}}$. Second, by our assumption there exists $\varepsilon > 0$ such that $\frac{\partial^2 v}{\partial\theta_f \partial q} \geqslant \varepsilon$. Then $\phi'_{\overline{\theta},\underline{\theta}}$ is positive and uniformly bounded away from zero

$$\phi'_{\overline{\theta},\underline{\theta}}(q) = v'_{\overline{\theta}}(q) - v'_{\underline{\theta}}(q) = \int_{\underline{\theta}}^{\overline{\theta}} \frac{\partial^2 v_{\theta_f}(q)}{\partial\theta_f \partial q} \mathrm{d}\theta_f \geqslant (\overline{\theta} - \underline{\theta})\varepsilon > 0.$$

□

Corollary 2 For any two of prices $\overline{p} \geqslant \underline{p} \geqslant 0$ and payoff types $\overline{\theta} > \underline{\theta}$, there exists q such that $u_{\underline{\theta}}(\underline{p},q) = u_{\overline{\theta}}(\overline{p},q)$.

Proof. Because $\phi_{\overline{\theta},\underline{\theta}}$ is invertible by Lemma 3, one can set $q := \phi^{-1}_{\overline{\theta},\underline{\theta}}(\overline{p} - \underline{p})$ so that it is as desired. □

Corollary 3 For any two payoff types $\overline{\theta} \geqslant \underline{\theta}$ and any contract $(p,q) \in \mathbb{R}^2$, there exist unique contracts $(\underline{p},\underline{q}),(\overline{p},\overline{q}) \in \mathbb{R}^2$, such that $\underline{q} \leqslant q \leqslant \overline{q}$ and

$$(p,q) \sim_{\underline{\theta}} (\underline{p},\underline{q}) \sim_{\overline{\theta}} \mathbf{0} \sim_{\underline{\theta}} (\overline{p},\overline{q}) \sim_{\overline{\theta}} (p,q).$$

Proof. Because $\phi_{\overline{\theta},\underline{\theta}}$ is invertible by Lemma 3, one can set

$$\underline{q} := \phi^{-1}_{\underline{\theta},\overline{\theta}}\big(u_{\underline{\theta}}(p,q)\big), \quad \underline{p} := v_{\overline{\theta}}(\underline{q}),$$

$$\overline{q} := \phi^{-1}_{\overline{\theta},\underline{\theta}}\big(u_{\overline{\theta}}(p,q)\big), \quad \overline{p} := v_{\underline{\theta}}(\overline{q}).$$

It is then straightforward to verify that $(\underline{p},\underline{q})$ and $(\overline{p},\overline{q})$ are as desired. □

Lemma 4 The efficient quality for payoff type θ_f defined as $\widehat{q}_{\theta_f} := \operatorname{argmax}_{q \geqslant 0} v_{\theta_f}(q) - \kappa(q)$ exists, is unique and increasing in θ_f.

Proof. For any payoff type $\theta_f \in \mathbb{R}$, define the corresponding surplus function $\zeta_{\theta_f} : \mathbb{R}_+ \to \mathbb{R}$ as

$$\zeta_{\theta_f}(q) := v_{\theta_f}(q) - \kappa(q) \tag{A.1}$$

and note our assumptions on v and κ imply that ζ is continuous, twice differentiable, initially strictly increasing ($\zeta'_{\theta_f}(0) > 0$), eventually strictly decreasing ($\lim_{q\to\infty} \zeta'_{\theta_f}(q) < 0$), strictly concave, and has increasing marginal differences ($\partial^2 \zeta_{\theta_f}(q)/\partial\theta_f \partial q \geqslant 0$). Therefore, the efficient quantity as defined in (9) is well-defined, is the unique maximizer of the surplus, and is strictly increasing in θ_f by standard monotone comparative statics arguments (see e.g. Edlin and Shannon, 1998). □

A.2. *Proofs*

Proof of Observation 1 on page 11: As a 1-EDP uses only a single frame, (i) implies (ii). By backward induction, an EDP using only a single frame is equivalent to a menu in the same frame comprising of all options offered somewhere in the EDP. Therefore, a $\mathbf{c}$ implemented by such an EDP needs to satisfy the $\mathrm{IC}^f_{\theta\theta'}$ and P^f_θ constraints and (ii) implies (iii). If $\mathbf{c}$ satisfies these constraints, the 1-EDP $(\{c_\theta\}_{\theta\in\Theta} \cup \{\mathbf{0}\}, f)$ implements $\mathbf{c}$, whence (iii) implies (i). □

We prove the two constitutive propositions before proceeding to the proof of Theorem 1.

Proof of Proposition 1 on page 14: The necessity of the constraints for a given set of frames $f_R, \{f_\theta\}_{\theta\in\Theta^C}$ is derived in the main text. In particular, we saw that the incentive compatibility constraints are determined by the frames used on the path to the contract of the imitated type, not the imitating type and frames used on the path to $c_{\theta'}$ cannot eliminate the IC constraints from θ' to θ for any $\theta', \theta \in \Theta$.

To prove the proposition, it remains to show that we can assume that $f_R = h$ and $f_\theta = \ell$. Suppose towards a contradiction that there exists a nonnegative $\mathbf{c}$ that is implementable by an EDP but does not satisfy the constraints of the proposition, and fix an f_R, Θ_C, and $\{f_\theta\}_{\theta\in\Theta_C}$ such that the associated constraints are satisfied, which must exist.

First, suppose $f_R \neq h$. But then $\mathbf{c}$ satisfies the constraints of the proposition with $\Theta'_C = \Theta$. As all contracts satisfy the participation constraint in $f_R < h$ and $q^\theta \geqslant 0$ for all $\theta \in \Theta$, they satisfy the participation constraint in ℓ by single crossing. As there are no incentive compatibility constraints with $\Theta'_C = \Theta$, all constraints associated to this set of hidden types are satisfied.

Second, suppose instead that $f_R = h$ but for some type $\theta' \in \Theta_C$ we have $f_{\theta'} \neq \ell$. But then the participation constraint for $f_{\theta'} = \ell$ is satisfied by single-crossing since $q^\theta \geqslant 0$. Hence, the set of contracts is feasible with $f_R = h$ and $f_\theta = \ell$ for the same Θ_C. □

Before we proceed, we prove two more detailed decoy construction lemmas, one for upward and an analogous one for downward deviations.

Lemma 5 For any increasing type profile $(\theta^i)_{i=0}^n$ and a nonnegative contract $d_0=(p_0,q_0)$, there exists a vector of decoy contracts $\mathbf{d}=(d_i)_{i=1}^n=(p_i,q_i)_{i=1}^n$, such that

(i) quantities are increasing: for $i\in\{1,\dots,n\}$, $q_i\geqslant q_{i-1}$;
(ii) all decoys are IC in frame h: for $i,j\in\{0,\dots,n\}$, $d_i\succsim_i^h d_j$;
(iii) decoy contracts are undesirable in frame ℓ: for $i\in\{1,\dots,n\}$, $\mathbf{0}\succsim_i^\ell d_i$.

Proof. The decoy contracts are constructed iteratively. For $i\in\{1,\dots,n\}$, obtain $(p_i,q_i):=(\overline{p},\overline{q})$ from Corollary 3 for $(p,q):=(p_{i-1},q_{i-1}),\underline{\theta}:=\theta_\ell^i,\overline{\theta}:=\theta_h^i$. Now note that (i) and (iii) follow immediately from Corollary 3, (ii) follows from Corollary 3 and the single crossing property (Lemma 2). □

Lemma 6 For any increasing type profile $(\theta^i)_{i=1}^{n+1}$ and a nonnegative contract $d_{n+1}=(p_{n+1},q_{n+1})\succsim_{\theta_l^{n+1}}\mathbf{0}$, either $d_{n+1}\precsim_i^\ell\mathbf{0}$ for all $i\in\{1,\dots,n\}$, or there exists a vector of decoy contracts $\mathbf{d}=(d_i)_{i=1}^n=(p_i,q_i)_{i=1}^n$, such that

(i) quantities are increasing: for $i\in\{1,\dots,n\}$, $q_i\geqslant q_{i-1}$;
(ii) all contracts are IC in frame ℓ: for $i,j\in\{0,\dots,n\}$, $d_i\succsim_i^\ell d_j$;
(iii) decoy contracts are undesirable in frame h: for $i\in\{1,\dots,n\}$, $\mathbf{0}\succsim_i^h d_i$.

Proof. The decoy contracts are constructed iteratively. For $i\in\{1,\dots,n\}$, obtain $(p_i,q_i):=(\underline{p},\underline{q})$ from Corollary 3 for $\underline{\theta}:=\theta^{i_l},\overline{\theta}:=\theta_h^i$, and

$$(p,q):=\begin{cases}(p_{i+1},q_{i+1}), & (p_{i+1},q_{i+1})\succsim_i^\ell\mathbf{0},\\ \mathbf{0}, & (p_{i+1},q_{i+1})\prec_i^\ell\mathbf{0}.\end{cases}$$

Now note that (i) and (iii) follow immediately from Corollary 3, (ii) follows from Corollary 3 and the single crossing property (Lemma 2). □

Proof of Lemma 1 on page 14: A continuation problem for type $\theta\in\Theta$ with a nonnegative contract c_θ satisfying all three properties is given by $e_\theta=\big(\{\mathbf{0},(\{\mathbf{0},c_\theta\}\cup\{d_{\theta'}\}_{\theta'>\theta},h)\},\cup\{d_{\theta'}\}_{\theta'<\theta},h)\},\ell\big)$, where the contracts $(d_{\theta'})_{\theta'>\theta}$ are constructed as $\mathbf{d}$ in Lemma 5 for the contract $d_0=c_\theta$, and type profile $(\theta')_{\theta'\geqslant\theta}$ and the contracts $(d_{\theta'})_{\theta'<\theta}$ are constructed as $\mathbf{d}$ in Lemma 6 for the contract $d_{n+1}=c_\theta$ and type profile $(\theta')_{\theta'\leqslant\theta}$.

By construction, type θ chooses c_θ from the terminal problem and since c_θ satisfies the participation constraint in the low frame, $c_\theta\in\Sigma^{e_\theta}(\theta)$. For higher types, the terminal decision problem resolves to the menu $\{d_{\theta'},0\}$ and by construction the outside option is weakly preferred in the low frame. Consider a type $\theta'<\theta$. In the terminal decision problem, we have $d_0\succsim_{\theta_h}d_i$ and $q_i\geqslant q_0$, hence by single crossing $d_0\succ_{\theta_h'}d_i$, which establishes that a lower type never chooses any of the decoys at the terminal stage. By construction, $d_{\theta'}\in\Sigma^{e_\theta}(\theta')$ and $d_{\theta'}\precsim_{\theta'}\mathbf{0}$ □

Proof of Proposition 2 on page 14: Suppose that a non-negative $\mathbf{c}$ satisfies $\{\mathrm{P}_\theta^h\}_{\theta\in\Theta_R}$, $\{\mathrm{P}_\theta^\ell\}_{\theta\in\Theta_C}$, $\{\mathrm{IC}_{\theta\theta'}^h\}_{\theta\in\Theta,\theta'\in\Theta_R}$ for some partition $\{\Theta_C,\Theta_R\}$ of Θ. Then let e_θ for each $\theta\in\Theta_C$ be constructed as in Lemma 1 and consider a canonical EDP

$$e^*=\Big(\{e_\theta\}_{\theta\in\Theta_C}\cup\{c_\theta\}_{\theta\in\Theta_R}\cup\{\mathbf{0}\},h\Big).$$

Notice that from Lemma 1 it follows for each $\theta\in\Theta_C$ that there exist a solution σ^θ of e_θ, such that

$$\sigma^\theta(\theta')\precsim_{\theta_h}c_\theta,\forall\theta'\in\Theta,$$

$$\sigma^\theta(\theta)=c_\theta.$$

Now let σ be such that $\sigma(\theta)=c_\theta$. To show that σ is a solution of e^*, notice that constraints $\{\mathrm{IC}_{\theta\theta'}^h\}_{\theta\in\Theta,\theta'\in\Theta_R}$ and $\{\mathrm{P}_\theta^h\}_{\theta\in\Theta_R}$ imply that $\forall\theta\in\Theta$,

$$\sigma(\theta)=c_\theta\in\operatorname*{argmax}_{\{c_{\theta'}\}_{\theta'\in\Theta_R}}u_{\theta_h}.$$

Therefore, σ satisfies

$$\sigma(\theta)\in\operatorname*{argmax}_{\{\sigma^{\theta'}(\theta)\}_{e_\theta}\cup\{c_{\theta'}\}_{\theta'\in\Theta_R}}u_{\theta_h},$$

which means that it is a solution of e^*.

Finally, let Q denote the set of outcomes in all contracts available in e^*. Then, setting $\underline{q}:=\min_i Q,\overline{q}:=\max_i Q$ ensures e^* is well-defined for any set of outcomes $\mathcal{Q}\supset[\underline{q},\overline{q}]$. □

Proof of Theorem 1 on page 12: Fix any nonnegative contract **c**. Then, for any $\mathcal{Q}$, (ii) always trivially implies (i), and (i) implies (iii) by Proposition 1. Finally, by Proposition 2, (iii) implies (ii) for a sufficiently large space of outcomes $\mathcal{Q}$. □

Proof of Theorem 2 on page 20: First, let us establish that for any Θ_C, a solution to the interior optimization problem in (RP) exists. Since $\lim_{q\to\infty} v'_{\theta_f}(q) - \kappa'(q) < 0$, the set of q that yield a nonnegative surplus is compact and the surplus attainable for any type and frame is bounded. We can therefore restrict the choice of contracts to those that have non-negative total surplus and transfers that do not exceed this bound: If the first condition is violated, the principal could increase profits by instead offering the outside option to all such types. If the second condition is violated, the participation constraint has to be violated. As the restricted set of contracts is compact and the objective and constraints are continuous, we have a solution.

Let $((c_\theta)_{\theta\in\Theta}, \Theta_C)$ be a solution to the relaxed problem and e^* a canonical EDP with decoys for all $\theta \in \Theta_C$ constructed as in Lemma 5 but omitting decoys for lower types as constructed in Lemma 6. Note that this EDP is feasible as all contracts and decoys have nonnegative quantities. We need to show that e^* implements $(c_\theta)_{\theta\in\Theta}$.

First, note that Σ^e is rectangular, i.e. if $\sigma, \sigma' \in \Sigma^e$ with $\sigma(\theta) \neq \sigma'(\theta)$ and $\sigma(\theta') \neq \sigma'(\theta')$, there exists a $\sigma^* \in \Sigma^e$ with $\sigma^* = \sigma$ except $\sigma^*(\theta') = \sigma'(\theta')$.

It follows from the IC constraints that there is no strictly profitable deviation into contracts of revealed types, i.e. $\Sigma^{e^*}(\theta) \cap \{c_{\theta'}\}_{\theta' \in \Theta_R \setminus \theta} \neq \emptyset$ implies that $c_\theta \in \Sigma^{e^*}(\theta)$. From Lemma 1, it follows that type θ cannot deviate downwards into concealed types and that no decoys are chosen, i.e. $\Sigma^{e^*}(\theta) \subset \{c_{\theta'}\}_{\Theta \setminus \{\theta' < \theta:\, \theta' \in \Theta_C\}}$. It remains to show that there are no strictly profitable upwards deviations in e^* to complete the proof, establishing $c_\theta \in \Sigma^{e^*}(\theta)$ for all $(c_\theta)_{\theta\in\Theta}$.

As the proof relies on properties of the solution to (RP), we start by simplifying the relaxed problem. Define

$$\beta(\theta) := \max\left\{\theta' \in \Theta_R :\, \theta' < \theta\right\}$$

the closest revealed type below a given type θ, and

$$\alpha(\theta) := \min\left\{\theta' \in \Theta_R :\, \theta' > \theta\right\}$$

the closest revealed type above a given type θ, with the notational convention that $\max\emptyset = \min\emptyset = \emptyset$. We now define the doubly relaxed problem, where we remove all but the downward IC constraints into the closest revealed type and the upwards IC constraints into the next largest revealed type.

$$\max_{\Theta_C \subseteq \Theta} \max_{\{(p_\theta, q_\theta)\}_{\theta\in\Theta}} \sum_{\theta\in\Theta} \mu_\theta\,(p_\theta - \kappa(q_\theta)) \tag{DRP}$$

$$\begin{aligned}
\text{s.t. } & v_{\theta_h}(q_\theta) - p_\theta \geqslant 0 \quad \forall \theta \in \Theta_R \\
& v_{\theta_\ell}(q_\theta) - p_\theta \geqslant 0 \quad \forall \theta \in \Theta_C \\
& v_{\theta_h}(q_\theta) - p_\theta \geqslant v_{\theta_h}(q_{\beta(\theta)}) - p_{\beta(\theta)} \quad \forall \theta \in \Theta \\
& v_{\theta_h}(q_\theta) - p_\theta \geqslant v_{\theta_h}(q_{\alpha(\theta)}) - p_{\alpha(\theta)} \quad \forall \theta \in \Theta.
\end{aligned}$$

We have the following

Lemma 7 The solution to (DRP) solves (RP) and satisfies R-monotonicity:

$$\theta, \theta' \in \Theta_R, \theta > \theta' \implies q_\theta \geqslant q_{\theta'}.$$

Proof. Note that transitivity implies it is sufficient to establish R-monotonicity only for adjacent revealed types. Take any $\theta, \theta' \in \Theta_R, \theta > \theta'$ and $\theta = \alpha(\theta'), \theta' = \beta(\theta)$. Then, the IC constraints

$$v_{\theta_h}(q_\theta) - p_\theta \geqslant v_{\theta_h}(q_{\beta(\theta)}) - p_{\beta(\theta)} = v_{\theta'_h}(q_{\theta'}) - p_{\theta'}$$

$$v_{\theta'_h}(q_{\theta'}) - p_{\theta'} \geqslant v_{\theta'_h}(q_{\alpha(\theta')}) - p_{\alpha(\theta')} = v_{\theta'_h}(q_\theta) - p_\theta$$

imply

$$v_{\theta_h}(q_\theta) - v_{\theta'_h}(q_\theta) \geqslant v_{\theta_h}(q_\theta) - v_{\theta'_h}(q_\theta).$$

Finally, $q_\theta \geqslant q_{\theta'}$ because the function $q \mapsto v_{\theta_h}(q) - v_{\theta'_h}(q) = \int_{\theta'_h}^{\theta_h} \frac{\partial v(t,q)}{\partial t}\,\mathrm{d}t$ is increasing.

To show that the solution to (DRP) solves (RP), it suffices to show that local IC imply global IC. Let us proceed by induction on the number of types in Θ_R between the source of the $\mathrm{IC}^h_{\theta\theta'}$ constraint θ and it's target θ'. If there are no revealed types between, then $\theta' = \beta(\theta)$ (resp. $\alpha(\theta)$) and we are done. Suppose that all constraints with up to n intermediate revealed types are implied and let $\theta' > \theta$, $\theta' \in \Theta_R$ with $n+1$ intermediate revealed types (the $\theta > \theta'$ case is identical). Then

$$\begin{aligned} v_{\theta_h}(q_\theta) - p_\theta &\geqslant v_{\theta_h}(q_{\beta(\theta)}) - p_{\beta(\theta)} \\ &= \big(v_{\theta_h}(q_{\beta(\theta)}) - v(\beta_h(\theta), q_{\beta(\theta)})\big) + v(\beta_h(\theta), q_{\beta(\theta)}) - p_{\beta(\theta)} \\ &\geqslant \big(v_{\theta_h}(q_{\beta(\theta)}) - v(\beta_h(\theta), q_{\beta(\theta)})\big) + v(\beta_h(\theta), q_{\theta'}) - p_{\theta'} \\ &\geqslant \big(v_{\theta_h}(q_{\theta'}) - v(\beta_h(\theta), q_{\theta'})\big) + v(\beta_h(\theta), q_{\theta'}) - p_{\theta'} \\ &= v_{\theta_h}(q_{\theta'}) - p_{\theta'} \end{aligned}$$

where we used the local IC, the induction hypothesis and monotonicity. Hence, all constraints of (RP) are satisfied by the solution to (DRP). □

Lemma 8 In the optimal contract of the relaxed problem the IC from any revealed type θ to the closest lower revealed type $\beta(\theta)$ is active.

Proof. As the relaxed problem and the doubly relaxed problem are equivalent, it is sufficient to show that local downward IC between revealed types are active in the doubly relaxed problem. Suppose towards a contradiction that one of them is not active, say from type θ to $\beta(\theta)$. Suppose we increase the price in the contract of all revealed types greater than θ including θ by some $\varepsilon > 0$. Note that this change does not affect any constraints between the affected types. Furthermore, θ is not the lowest revealed type, hence the participation constraint of all revealed type is implied by the IC and not active since IC-$\theta \to \beta(\theta)$) is not active. As we can pick epsilon sufficiently small, this IC is still slack and we strictly increased revenue, contradiction the optimality of the initial contract. □

Lemma 9 In the optimal contract of the relaxed problem, $q_\theta \leqslant \widehat{q}_{\theta_h}$ for all $\theta \in \Theta$.

Proof. As the relaxed problem and the doubly relaxed problem are equivalent, we can work on the doubly relaxed problem. The result follows from Lemma 4 for concealed types. Suppose towards a contradiction that this property is violated for some subset of revealed types. Pick the smallest revealed type for which this is the case and denote it as θ. Note that $q_{\beta(\theta)} \leqslant \widehat{q}_{\beta_h(\theta)} < \widehat{q}_{\theta_h} < q_{\theta_h}$ and denote the rent given to type θ as $\Delta := v(\theta_h, q_{\beta(\theta)}) - p_{\beta(\theta)}$. (This is the correct expression, because the local downward IC is active by the above lemma.) Consider the set of contracts where we replaced the initial contract for type θ by $\big(\widehat{q}_{\theta_h}, v_{\theta_h}(\widehat{q}_{\theta_h}) - \Delta\big)$. As θ receives the same utility in both contracts, no participation constraint is violated and all IC from θ are still satisfied. The upward IC $\beta(\theta) \to \theta$ is still satisfied as it is implied by R-monotonicity (which is maintained) and the corresponding downward IC. Consider any higher type imitating θ. The amended contract gives the same utility to θ at a lower quality, hence it gives a strictly lower deviation payoff to higher types. In particular, all IC are satisfied. The revised contract is also more profitable for the principal as the most profitable way to transfer rent to type θ in frame h is using quality $\widehat{q}_{\theta_h}$. Hence, the initial set of contracts was not optimal. □

Now, we can show that there are no profitable feasible upward deviations in e^*. We proceed by induction. Order the types such that $\{\theta^1, \ldots, \theta^n\} = \Theta$, $\theta^i < \theta^{i+1}$. Clearly, the highest type has no feasible upward deviations. Suppose all upward deviations are either infeasible or unprofitable for types θ^i into types θ^j for $j > i > m$. We need to show that the required upward IC constraints out of type θ^m are satisfied. We will proceed case by case, in addition showing that the upward IC from concealed to revealed types are always slack:

1. *Deviations into a concealed type with rent* $\Delta_{\theta^i} \leqslant v(\theta^i_h, \widehat{q}_{\theta^i_h}) - v(\theta^i_\ell, \widehat{q}_{\theta^i_h})$*:* Then, the participation constraint of type θ^i is binding at the intermediate stage in frame ℓ. But by single crossing

$$c_{\theta^i} \sim_{\theta^i_l} \mathbf{0} \Longrightarrow c_{\theta^i} \prec_{\theta^m_l} \mathbf{0}, \tag{A.2}$$

an imitation is infeasible.

2. *Deviations into a concealed type with rent* $\Delta_{\theta^i} > v(\theta^i_h, \widehat{q}_{\theta^i_h}) - v(\theta^i_\ell, \widehat{q}_{\theta^i_h})$*:* Note that in this case $q_{\theta^i} = \widehat{q}_{\theta^i_h}$ and this rent has to be the result of a possible deviation that is discouraged by a constraint of the problem and hence by the induction hypothesis this is a downward deviation into a revealed type. Hence, $\Delta_{\theta^i} = v(\theta^i_h, q_\eta) - p_\eta$ for some $\eta < \theta^i$, $\eta \in \Theta_R$. But then the upward deviation is not profitable unless the deviation into η is profitable, since $q_\eta \leqslant \widehat{q}_{\eta_h} < \widehat{q}_{\theta^i_h} = q_{\theta^i}$ and by single crossing

$$c_\eta \sim_{\theta^i_h} c_{\theta^i} \Longrightarrow c_\eta \succ_{\theta^m_h} c_{\theta^i} \tag{A.3}$$

so all we have to show is that deviations into revealed types are not profitable. If $\eta < \theta^m$, this is achieved already by the maintained IC constraints, if $\theta^m \in \Theta_R$ it is by the upward IC. The case we need to consider are deviations from concealed types upwards into revealed types.

3. *Deviations from a concealed into a revealed type:* Consider a concealed type θ^m with a profitable upwards deviation into a revealed type. As the set of types is finite, there has to exist a lowest revealed type into which θ^m has a strictly profitable deviation. Furthermore, since we impose downward incentive compatibility constraints, this lowest target type has to be greater than θ^m. We will show that such a lower bound cannot exist, hence there can be no profitable upward deviation.

Suppose such a lower bound exists, $\underline{\theta} = \min\{\theta \in \Theta_R : v_{\theta_h^m}(q_\theta) - p_\theta > v_{\theta_h^m}(q_{\theta^m}) - p_{\theta^m}\}$. But then, consider type $\beta(\underline{\theta})$. A deviation into this type is also strictly profitable since $c_{\beta(\underline{\theta})} \sim_{\underline{\theta}_h} c_{\underline{\theta}}$ and by R-monotonicity $q_{\beta(\underline{\theta})} \leqslant q_{\underline{\theta}}$, but then by single crossing $c_{\beta(\underline{\theta})} \succsim_{\theta_h^m} c_{\underline{\theta}} \succ_{\theta_h^m} c_{\theta^m}$, contradicting the minimality of $\underline{\theta}$. If $\beta(\underline{\theta}) = \emptyset$ an analogous argument establishes that the participation constraint is violated. Hence, there can be no strictly profitable upward deviation.

And we established that there can be no upward deviation by type θ^m. By induction, no type prefers any attainable contract offered to higher types in e^*, and hence, we found an EDP that attains the upper bound to the solution of (GP) and therefore (RP)=(GP). □

Proof of Proposition 3 on page 21: Let $\mathbf{c} = (c_\theta)_\theta = ((p^*_\theta, q^*_\theta))_\theta$ be an optimal vector of contracts implemented by some EDP. By Theorem 1, we can construct a canonical EDP e with that implements it. Let Θ_C and Θ_R be the sets of revealed and concealed types in e. If $\theta \in \Theta_C$, the statement follows from Lemma 4. We proved that $q^*_\theta < \widehat{q}_{\theta_h}$ as Lemma 9. Therefore, there is only one case left to consider. Assume that $\theta \in \Theta_R$ and towards a contradiction that $q^*_\theta < q_\theta$, where q_θ satisfies $\zeta_{\theta_h}(q_\theta) = \zeta_{\theta_\ell}(\widehat{q}_{\theta_\ell})$, where ζ_θ is the surplus function defined in (4). Denote the rent in this contract by $\Delta := v_{\theta_h}(q^*_\theta) - p^*_\theta$.

We will construct a vector of contracts with strictly higher revenue. Starting from the canonical EDP, we now conceal type θ and set the contract $(\widehat{q}_{\theta_\ell}, \widehat{p}_{\theta_\ell} - \Delta)$. Note that since $q^*_\theta < q_\theta$, we have $\zeta_{\theta_h}(q_\theta) < \zeta_{\theta_\ell}(\widehat{q}_{\theta_\ell})$ and consequently

$$\zeta_{\theta_h}(q_\theta) - \Delta < \zeta_{\theta_\ell}(\widehat{q}_{\theta_\ell}) - \Delta$$

$$p_\theta - \kappa(q_\theta) < \widehat{p}_{\theta_\ell} - \Delta - \kappa(\widehat{q}_{\theta_\ell})$$

and the principal receives weakly higher profit in the modified contract.

Clearly, this contract satisfies the participation constraint in frame ℓ and delivers rent greater than Δ to type θ in the high frame, hence there is no deviation by this type. There is no downward deviation into this contract since the type is concealed. Furthermore, we do not have to worry about upward deviations. The optimal concealed contract—which delivers even higher profits—is never subject to them and we have established that even a sub-optimal concealed contract delivers an improvement in profits. Hence, the original vector was not optimal, a contradiction. □

Proof of Lemma 4 on page 21: Note that there are no IC constraints into a type $\theta \in \Theta_C$. Hence, we can separate the principals problem and solve for the optimal contract of θ in (RP). The contract given to type θ solves

$$\max_{(p,q)} p - \kappa(q)$$

$$\text{s.t. } v_{\theta_\ell}(q) - p \geqslant 0$$

$$v_{\theta_h}(q) - p \geqslant \Delta$$

Dropping the second constraint, the optimal contract is $\widehat{c}_{\theta_\ell}$, which delivers rent $v_{\theta_h}(\widehat{q}_{\theta_\ell}) - v_{\theta_\ell}(\widehat{q}_{\theta_\ell})$, hence the second constraint is satisfied if

$$\Delta \leqslant \left[v_{\theta_h}(\widehat{q}_{\theta_\ell}) - v_{\theta_\ell}(\widehat{q}_{\theta_\ell}) \right]. \tag{A.4}$$

Similarly, note that the optimal contract dropping the first constraint is $\left(v_{\theta_h}(\widehat{q}_{\theta_h}) - \Delta, \widehat{q}_{\theta_h}\right)$, which gives utility $v_{\theta_\ell}(\widehat{q}_{\theta_h}) - v_{\theta_h}(\widehat{q}_{\theta_h}) + \Delta$ in the low frame. Hence, the first constraint is satisfied if

$$\Delta \geqslant v_{\theta_h}(\widehat{q}_{\theta_h}) - v_{\theta_\ell}(\widehat{q}_{\theta_h}). \tag{A.5}$$

In the intermediate case, both constraints are binding,

$$v_{\theta_\ell}(q^*) = p$$

$$v_{\theta_h}(q^*) - v_{\theta_\ell}(q^*) = \Delta$$

and the optimal contract is $\left(v_{\theta_\ell}(q^*), q^*\right)$. Note that $q^* \in (\widehat{q}_{\theta_\ell}, \widehat{q}_{\theta_h})$ by single crossing. □

Proof of Proposition 5 on page 23: Take any type $\theta \in \Theta$. For each μ, consider (RP) with the constraint $\theta \in \Theta_C$ $(\theta \in \Theta_R)$ and denote the corresponding optimal value by $\Pi^R_{C,\mu}$ $(\Pi^R_{R,\mu})$. Next, using the surplus function ζ_{θ_f} defined in (A.1), we can bound those values as

$$\Pi^R_{R,\mu} \geqslant \mu_\theta \zeta_{\theta_h}(\widehat{q}_{\theta_h}) \tag{A.6}$$

$$\Pi^R_{C,\mu} \leqslant \mu_\theta \zeta_{\theta_\ell}(\widehat{q}_{\theta_\ell}) + \sum_{\theta' \neq \theta} \mu_{\theta'} \zeta_{\theta'_h}(\widehat{q}_{\theta'_h}) \leqslant \mu_\theta \zeta_{\theta_\ell}(\widehat{q}_{\theta_\ell}) + (1 - \mu_\theta) \zeta_{\tilde{\theta}_h}(\widehat{q}_{\tilde{\theta}_h}), \tag{A.7}$$

where $\tilde{\theta} := \max(\Theta \setminus \{\theta\})$.

Note that Lemma 4 implies that

$$\zeta_{\theta_h}(\widehat{q}_{\theta_h}) > \zeta_{\theta_h}(\widehat{q}_{\theta_\ell}) > \zeta_{\theta_\ell}(\widehat{q}_{\theta_\ell}), \tag{A.8}$$

and define

$$\overline{\mu}_\theta := \frac{\zeta_{\tilde{\theta}_h}(\widehat{q}_{\tilde{\theta}_h})}{\zeta_{\theta_h}(\widehat{q}_{\theta_h}) - \zeta_{\theta_\ell}(\widehat{q}_{\theta_\ell}) + \zeta_{\tilde{\theta}_h}(\widehat{q}_{\tilde{\theta}_h})} \in (0,1).$$

Finally, combining (A.6), (A.7), and (A.8) yields

$$\begin{aligned}\Pi^R_{R,\mu} - \Pi^R_{C,\mu} &\geqslant \mu_\theta \zeta_{\theta_h}(\widehat{q}_{\theta_h}) - \mu_\theta \zeta_{\theta_\ell}(\widehat{q}_{\theta_\ell}) - (1-\mu_\theta)\zeta_{\tilde{\theta}_h}(\widehat{q}_{\tilde{\theta}_h}) \\ &= \mu_\theta \left[\zeta_{\theta_h}(\widehat{q}_{\theta_h}) - \zeta_{\theta_\ell}(\widehat{q}_{\theta_\ell}) + \zeta_{\tilde{\theta}_h}(\widehat{q}_{\tilde{\theta}_h})\right] - \zeta_{\tilde{\theta}_h}(\widehat{q}_{\tilde{\theta}_h}) \\ &\geqslant 0.\end{aligned}$$

Therefore, for any $\mu_\theta \in [\overline{\mu}_\theta, 1]$, it is optimal to reveal θ.[32] □

Proof of Proposition 6 on page 23: It is easy to see that the optimal contract and set of concealed types before the change of valuations is still feasible after the change. Hence $\Pi^*_\Theta \leqslant \Pi^*_{\tilde{\Theta}}$. If $\tilde{\theta}$ is concealed in the optimum, we are done. Suppose that instead it is not concealed. Then, since $\tilde{\theta}_l$ is the only difference to the initial problem and does not affect the constraints unless $\tilde{\theta}$ is concealed, the optimal contract under $\tilde{\Theta}$ is feasible under Θ and $\Pi^*_{\tilde{\Theta}} \leqslant \Pi^*_\Theta$. Hence, the original vector of contracts is still optimal, establishing the claim. □

Proof of Proposition 7 on page 23: Consider the profit from concealing all types except the highest, $\Pi_C(\varepsilon) := \sum_{\theta<\overline{\theta}} \mu_\theta v_{\theta_\ell}(\widehat{q}_{\theta_\ell}) + \mu_{\overline{\theta}} v_{\overline{\theta}_h}(\widehat{q}_{\overline{\theta}_h})$, where we consider the θ_ℓ as a function of ε. It is easy to see that $\widehat{q}_{\theta_\ell} \to \widehat{q}_{\theta_h}$ as $\theta_\ell \to \theta_h$. By continuity of v, $\Pi_C(\varepsilon) \to \Pi((\widehat{c}_{\theta_h})_{\theta\in\Theta})$.

Suppose that in the optimum $\theta < \overline{\theta}$ is revealed. Then $q_\theta > \underline{q_\theta} > 0$ by Proposition 3 and $p_\theta \leqslant v_{\theta_h}(q_\theta)$. But then, by the incentive compatibility constraint of $\overline{\theta}$, $\Pi < \Pi((\widehat{c}_{\theta_h})_{\theta\in\Theta}) - \mu_{\overline{\theta}}(v_{\overline{\theta}_h}(q_\theta) - v_{\theta_h}(q_\theta)) < \Pi((\widehat{c}_{\theta_h})_{\theta\in\Theta})$. Hence, there exists an $\varepsilon_\theta > 0$ such that $\Pi_C(\varepsilon) > \Pi$ for $\varepsilon < \varepsilon_\theta$, so it cannot have been optimal to reveal θ for sufficiently small ε. The result follows by taking the maximum over $\{\varepsilon_\theta : \theta \in \Theta \setminus \overline{\theta}\}$. □

Proof of Theorem 3 on page 25: Let e_0 denote an EDP constructed for sophisticated types in Theorem 1. Order naive types $\Theta_N = \{\theta^1, \ldots, \theta^m\}$ with $\theta^i < \theta^{i+1}$. We will construct an optimal EDP for the mixed case inductively.

Starting from $e_0 = (E_0, h)$, we add one continuation problem at the root for every naive type,

$$e_{n+1} = \left(\bigcup_{i=0}^{n+1} E_i, h\right). \tag{A.9}$$

To define E_i, let the most preferred alternative in e_{i-1} for type θ^i be $x_i := \operatorname{argmax}_{\mathcal{C}(e_{i-1})} u_{\theta^i_h}$. During the construction, we ensure that

1. no sophisticated type prefers to continue to E_i,
2. no naive type θ^j with $j < i$ prefers to continue to E_i, and
3. type θ^i indeed proceeds to E_i and chooses $\widehat{c}_{\theta^i_h}$ eventually.

If we ensure this during our construction, all sophisticated types choose as in e_0 and all naive types choose their efficient contract $\widehat{c}_{\theta^i_h}$ and we establish the theorem.

Let $E_i = \left\{\left(\left\{\left(\{N_i, \{d^{\theta'}_{N,i}\}_{\theta'>\theta^i:\,\theta'\in\Theta_S}, \mathbf{0}\}, h\right), \left(\{\widehat{c}_{\theta^i_h}, \{d^{\theta'}_i\}_{\theta'>\theta^i:\,\theta'\in\Theta_S}, \mathbf{0}\}, h\right), \mathbf{0}\right\}, \ell\right)\right\}$.

We now have to specify N_i and the decoys and verify 1–3 above. First, use Corollary 3 to construct $N_i = (\overline{p}, \overline{q})$ for $(p,q) := x_i, \overline{\theta} := \theta^i_h, \underline{\theta} := \theta^i_l$ so that

$$N_i \sim_{\theta^i_h} x_i \tag{A.10}$$

$$q_{N_i} \geqslant q_{x_i} \tag{A.11}$$

$$N_i \preccurlyeq_{\theta^i_l} \mathbf{0} \tag{A.12}$$

Second, let $\Theta_S^{>\theta^i}$ ($\Theta_S^{\geqslant\theta^i}$) be a vector of types in Θ_S that are (weakly) greater than θ^i and define the decoys $(d^{\theta'}_{N,i})_{\theta'\in\Theta_S^{>\theta^i}}$ as $\mathbf{d}$ from Lemma 5 for the contract $d_0 = N_i$ and type profile $\Theta_S^{\geqslant\theta^i}$, and decoys $(d^{\theta'}_i)_{\theta'\in\Theta_S^{>\theta^i}}$ as $\mathbf{d}$ from Lemma 5 for the contract $d_0 = \widehat{c}_{\theta^i_h}$ and type profile $\Theta_S^{\geqslant\theta^i}$.

32. This bound is typically not tight, as we introduced slack in (A.7).

By construction, every sophisticated type $\theta > \theta^i$ prefers the outside option to the contract chosen from the continuation problems. Hence, they have no incentive to enter. Furthermore, all contracts are, by construction, worse in frame ℓ than the outside option for all types $\theta < \theta^i$, hence lower sophisticated types have no incentive to enter. Hence, we established 1.

By construction, E_i contains a most preferred option for θ_h^i, hence continuing into E_i is part of a naive solution for θ^i. At the subsequent decision node, the decision problem containing N_i is as attractive as the outside option: By the construction of the decoys, $N_i \succsim_{\theta_h^i} d_{N,i}^{\theta'}$ and $q_{N_i} \leqslant q_{d_{N,i}^{\theta'}}$ and hence $N_i \succsim_{\theta_l^i} d_{N,i}^{\theta'}$. But $N_i \precsim_{\theta_l^i} \mathbf{0}$. As the decision problem containing $\widehat{c}_{\theta_h^i}$ also contains the outside option, continuing to this menu is part of a naive solution. From the menu $\left(\{\widehat{c}_{\theta_h^i}, \{d_i^{\theta'}\}_{\theta' > \theta^i:\, \theta' \in \Theta_S}, \mathbf{0}\}, h\right)$, the DM chooses $\widehat{c}_{\theta_h^i}$ by the construction of the decoys. This establishes 3.

To see 2, note that all decoys have higher quality than N_i and $\widehat{c}_{\theta_h^i}$, respectively, and are less preferred according to θ_h^i. Hence, they are less preferred by lower naive types θ_h^j by single crossing. Furthermore, $\widehat{c}_{\theta_h^i}$ is not attractive to lower naive types, as it is worse than the outside option. It remains to check whether N_i is attractive. But note that $N_i \sim_{\theta_h^i} x_i$ and $q_{N_i} \geqslant q_{x_i}$ imply $N_i \precsim_{\theta_h^j} x_i$ for all $j < i$. By the induction hypothesis, $N_j \succsim_{\theta_h^j} \operatorname{argmax}_{\mathcal{C}(e_{i-1})} u_{\theta_h^j} \succsim_{\theta_h^j} x_i \succsim_{\theta_h^j} N_i$. Consequently, N_i is not attractive to lower naive types, and there is a naive solution where types $\theta^j < \theta^i$ choose E_j.

Clearly, the contract implemented for naive types is optimal given the participation constraint in the high frame any implemented contract needs to satisfy. Furthermore, suppose there is an EDP implementing contracts for sophisticated types that are not implemented by an optimal EDP in Theorem 1. Then, the contracts do not solve (RP), so we can find a strictly better set of contracts and use the above construction. Hence, every optimal EDP in (11) satisfies Theorem 3. From that, the decomposition theorem is immediate. □

Proof of Observation 3 on page 26: Let us denote the contract for type θ in the sophisticated problem as c_θ^s and note that the contract in the naive problem is $\widehat{c}_{\theta_h}$. Note that $c_\theta^s \succsim_{\theta_h} \mathbf{0} \sim \widehat{c}_{\theta_h}$ and $q_\theta^s \leqslant \widehat{q}_{\theta_h}$. Hence by single crossing $c_\theta^s \succsim_{\theta_f} \widehat{c}_{\theta_h}$, strictly for $f \neq h$ if $c_\theta^s \neq \widehat{c}_{\theta_h}$. □

Proof of Observation 4 on page 27: To implement the vector of contracts $(\widehat{c}_{\theta_n})_{\theta\in\Theta}$, the principal can simply conceal *all* types using neutral frame n.

Notice that $(\widehat{c}_{\theta_n})_{\theta\in\Theta}$ satisfies all the constraints of (RP) for $\Theta_R = \varnothing, \Theta_C = \Theta$. Therefore, by Theorem 2, there exists a canonical EDP e^* that implements it. Notice that since the contract $\widehat{c}_{\theta_n}$ for type θ satisfy P_θ^n, the interim and *ex post* modifications $\overline{e}^*$ and $\underline{e}^*$ also implement $(\widehat{c}_{\theta_n})_{\theta\in\Theta}$. □

Proof of Observation 5 on page 28: First, consider the *ex post* modification. Any naive solution needs to satisfy $\nu(\theta) \succsim_{\theta_n} \mathbf{0}$. The revenue maximal vector of contracts satisfying these constraints is $(\widehat{c}_{\theta_n})_{\theta\in\Theta_N}$. It is immediate from the proof of Theorem 3 that this set of contracts can be implemented using an analogous construction.

Second, consider the interim modification. Suppose that the optimal EDP without the modification is e^* and notice that e^* implements $(\widehat{c}_{\theta_h})_{\theta\in\Theta}$. Now consider its interim modification $\overline{e}^*$. Since naive consumers think they would get a better option than 0, they would proceed to e^*. Therefore, there exists a naive solution $\overline{\nu}$ to $\overline{e}^*$, such that $\overline{\nu}_\theta = \widehat{c}_{\theta_h}$ for all θ. □

A.3. *Construction of Example 2*

It is apparent from Figure 4 that the procedure sketched out in the text reduces this deviation surplus to zero in the minimal number of steps.[33] Formally, let $d_n = (p_n, q_n)$ denote the sequence of decoy contracts numbered from the last stage of the problem towards the root, which we start with $d_0 = c_\eta$. Let

$$q_{2n+1} := \overline{q}, \quad p_{2n+1} := p_{2n} + \theta_h(\overline{q} - q_{2n})$$

for the decoys in odd periods, which are at the upper bound. The decoy in even periods satisfies

$$\theta_\ell q_{2n} - p_{2n} = \theta_\ell \overline{q} - p_{2n-1}$$
$$\eta_\ell q_{2n} - p_{2n} = \eta_\ell q_0 - p_0.$$

33. If c_θ satisfies the participation constraint in the low frame, it is possible to save one stage by starting in the low frame.

Solving for these contracts, we get[34]

$$q_{2n} = \frac{\theta_\ell \overline{q} - p_{2n-1} - (\eta_\ell q_0 - p_0)}{\theta_\ell - \eta_\ell}$$

$$p_{2n} = \frac{(\overline{q} - q_0)\theta_\ell \eta_\ell + p_0 \theta_\ell - \eta_\ell p_{2n-1}}{\theta_\ell - \eta_\ell}.$$

We can then substitute into the odd-index price and obtain the geometric series.

$$p_{2n+1} = \frac{(\theta_h - \eta_\ell) p_{2n-1} - \eta_\ell (\overline{q} - q_0)(\theta_h - \theta_\ell) + p_0 (\theta_h - \theta_\ell)}{\theta_\ell - \eta_\ell}.$$

Writing $a_n = p_{2n-1}$ we have $a_1 = p_0 + \theta_h(\overline{q} - q_0)$ and solving the recursion, we get

$$a_n = p_0 + (\overline{q} - q_0)\left(\eta_\ell + \left(\frac{\theta_h - \eta_\ell}{\theta_\ell - \eta_\ell}\right)^n (\theta_\ell - \eta_\ell)\right).$$

Recall that we have reduced the deviation surplus to zero when the price of the odd-index decoy (with quality $\overline{q}$) exceeds the low-frame willingness to pay $\theta_h \overline{q}$. We solve

$$\theta_\ell \overline{q} = p_0 + (\overline{q} - q_0)\left(\eta_\ell + \left(\frac{\theta_h - \eta_\ell}{\theta_\ell - \eta_\ell}\right)^n (\theta_\ell - \eta_\ell)\right)$$

$$n = \frac{\log\left(\frac{\overline{q}}{\overline{q} - q_0} + \frac{q_0 \eta_\ell - p_0}{(\overline{q} - q_0)(\theta_\ell - \eta_\ell)}\right)}{\log\left(1 + \frac{\theta_h - \theta_\ell}{\theta_\ell - \eta_\ell}\right)}.$$

Let $\lceil n \rceil$ denote the next larger integer. The procedure needs $2\lceil n \rceil + 1$ steps to implement the contract if we have to start in the high frame, and $2\lceil n \rceil$ if we can start in the low frame. The comparative statics mentioned in the text follow from straightforward computation.

REFERENCES

ARENI, C. S. and KIM, D. (1993), "The Influence of Background Music on Shopping Behavior: Classical versus Top-Forty Music in a Wine Store", in McAlister, L. and Rothschild, M. L. (eds) *ACR North American Advances* (Provo, UT: Association for Consumer Research) 336–340.

BERNHEIM, B. D. and RANGEL, A. (2009), "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics", *The Quarterly Journal of Economics*, **124**, 51–104.

BORDALO, P., GENNAIOLI, N. and SHLEIFER, A. (2013), "Salience and Consumer Choice", *Journal of Political Economy*, **121**, 803–843.

BUSHONG, B., KING, L. M., CAMERER, C. F. and RANGEL, A. (2010), "Pavlovian Processes in Consumer Choice: The Physical Presence of a Good Increases Willingness-to-Pay", *American Economic Review*, **100**, 1556–1571.

DE CLIPPEL, G. (2014), "Behavioral Implementation", *American Economic Review*, **104**, 2975–3002.

DELVECCHIO, D., SHANKER KRISHNAN, H. and SMITH, D. C. (2007), "Cents or Percent? The Effects of Promotion Framing on Price Expectations and Choice", *Journal of Marketing*, **71**, 158–170.

EDLIN, A. S. and SHANNON, C. (1998), "Strict Monotonicity in Comparative Statics", *Journal of Economic Theory*, **81**, 201–219.

ELIAZ, K. and SPIEGLER, R. (2006), "Contracting with Diversely Naive Agents", *The Review of Economic Studies*, **73**, 689–714.

34. Clearly the q_{2n} are decreasing. To guarantee that this quantity is always positive until we have reduced the surplus to zero, we need to ensure that the worst case penultimate decoy has a positive quantity. This point solves

$$q\theta_h - p = \overline{q}\theta_h - P$$

$$q\eta_\ell - p = q_0 \eta_\ell - p_0$$

$$\overline{q}\theta_\ell - P = 0$$

or equivalently $q = \frac{\overline{q}(\theta_h - \theta_\ell) - (q_0 \eta_\ell - p_0)}{\theta_h - \eta_\ell}$. Hence, we do not have to worry about the lower bound in constructing our decoys if $\overline{q}(\theta_h - \theta_\ell) - (q_0 \eta_\ell - p_0) \geqslant 0$. Otherwise, the intersection with the indifference curve of η_ℓ is replaced with the condition $q_{2n} = 0$ once the non-negativity constraint is binding, The rest of the construction operates unchanged. We assume that the condition is satisfied in the above to simplify the exposition.

——— and ———, (2008), "Consumer Optimism and Price Discrimination", *Theoretical Economics*, **3**, 459–497.

ESŐ, P. and SZENTES, B. (2007), "Optimal Information Disclosure in Auctions and the Handicap Auction", *The Review of Economic Studies*, **74**, 705–731.

ESTEBAN, S. and MIYAGAWA, E. (2006), "Optimal Menu of Menus with Self-Control Preferences" (mimeo).

——— and ———, (2006), "Temptation, Self-Control, and Competitive Nonlinear Pricing", *Economics Letters*, **90**, 348–355.

———, ———, and SHUM, M. (2007), "Nonlinear Pricing with Self-Control Preferences", *Journal of Economic Theory*, **135**, 306–338.

FREEMAN, D. J. (2021), "Revealing Naïveté and Sophistication from Procrastination and Preproperation", *American Economic Journal: Microeconomics*, **13**, 402–438.

GALPERTI, S. (2015), "Commitment, Flexibility, and Optimal Screening of Time Inconsistency", *Econometrica*, **83**, 1425–1465.

GLAZER, J. and RUBINSTEIN, A. (2012), "A Model of Persuasion with Boundedly Rational Agents", *Journal of Political Economy*, **120**, 1057–1082.

——— AND ———, (2014), "Complex Questionnaires", *Econometrica*, **82**, 1529–1541.

GOTTLIEB, D. and MITCHELL, O. S. (2020), "Narrow Framing and Long-Term Care Insurance", *Journal of Risk and Insurance*, **87**, 861–893.

——— AND ZHANG, X. (2021), "Long-Term Contracting With Time-Inconsistent Agents", *Econometrica*, **89**, 793–824.

GOURVILLE, J. T. (1998), "Pennies-a-Day: The Effect of Temporal Reframing on Transaction Evaluation", *Journal of Consumer Research*, **24**, 395–403.

GUL, F. and PESENDORFER, W. (2001), "Temptation and Self-Control", *Econometrica*, **69**, 1403–1435.

HARDISTY, D. J., JOHNSON, E. J. and WEBER, E. U. (2010), "A Dirty Word or a Dirty World? Attribute Framing, Political Affiliation, and Query Theory", *Psychological Science*, **21**, 86–92.

HEIDHUES, P. and KŐSZEGI, B. (2010), "Exploiting Naïvete about Self-Control in the Credit Market", *American Economic Review*, **100**, 2279–2303.

——— AND ———, (2017), "Naivete-Based Discrimination", *The Quarterly Journal of Economics*, **132**, 1019–1054.

HERRERO, M. J. and SRIVASTAVA, S. (1992), "Implementation via Backward Induction", *Journal of Economic Theory*, **56**, 70–88.

JOHNSON, E. J., HERSHEY, J., MESZAROS, J. and KUNREUTHER, H. (1993), "Framing, Probability Distortions, and Insurance Decisions", *Journal of Risk and Uncertainty*, **7**, 35–51.

KŐSZEGI, B. (2014), "Behavioral Contract Theory", *Journal of Economic Literature*, **52**, 1075–1118.

——— AND SZEIDL, A. (2013), "A Model of Focusing in Economic Choice", *The Quarterly Journal of Economics*, **128**, 53–104.

LAIBSON, D. (1997), "Golden Eggs and Hyperbolic Discounting", *The Quarterly Journal of Economics*, **112**, 443–478.

LI, H. and SHI, X. (2017), "Discriminatory Information Disclosure", *American Economic Review*, **107**, 3363–3385.

LOEWENSTEIN, G., O'DONOGHUE, T. and RABIN, M. (2003), "Projection Bias in Predicting Future Utility", *The Quarterly Journal of Economics*, **118**, 1209–1248.

MASKIN, E. and RILEY, J. (1984), "Monopoly with Incomplete Information", *The RAND Journal of Economics*, **15**, 171–196.

MICHEL, C. and STENZEL, A. (2021), "Model-based Evaluation of Cooling-off Policies", *Games and Economic Behavior*, **129**, 270–293.

MOORE, J. and REPULLO, R. (1988), "Subgame Perfect Implementation", *Econometrica*, **56**, 1191.

MOSER, C. and OLEA DE SOUZA E SILVA, P. (2019), "Optimal Paternalistic Savings Policies" (Working Paper).

MUSSA, M. and ROSEN, S. (1978), "Monopoly and Product Quality", *Journal of Economic Theory*, **18**, 301–317.

NORTH, A. C., SHILCOCK, A. and HARGREAVES, D. J. (2003), "The Effect of Musical Style on Restaurant Customers' Spending", *Environment and Behavior*, **35**, 712–718.

———, HARGREAVES, D. J. and MCKENDRICK, J. (1997), "In-Store Music Affects Product Choice", *Nature*, **390**, 132–132.

———, SHERIDAN, L. P. and ARENI, C. S. (2016), "Music Congruity Effects on Product Memory, Perception, and Choice", *Journal of Retailing*, **92**, 83–95.

PICCIONE, M. and SPIEGLER, R., (2012), "Price Competition Under Limited Comparability", *The Quarterly Journal of Economics*, **127**, 97–135.

SALANT, Y. and RUBINSTEIN, A. (2008), "(A, f): Choice with Frames", *The Review of Economic Studies*, **75**, 1287–1296.

——— AND SIEGEL, R. (2018), "Contracts with Framing", *American Economic Journal: Microeconomics*, **10**, 315–346.

SCHKADE, D. A. and KAHNEMAN, D. (1998), "Does Living in California Make People Happy? A Focusing Illusion in Judgments of Life Satisfaction", *Psychological Science*, **9**, 340–346.

SCHMITZ, H. and ZIEBARTH, N. R. (2017), "Does Price Framing Affect the Consumer Price Sensitivity of Health Plan Choice?", *Journal of Human Resources*, **52**, 88–127.

SCHWITZGEBEL, E. and CUSHMAN, F. (2015), "Philosophers' Biased Judgments Persist despite Training, Expertise and Reflection", *Cognition*, 2015, **141**, 127–137.

SPIEGLER, R. (2011), *Bounded Rationality and Industrial Organization* (Oxford University Press).

———, (2014), "Competitive Framing", *American Economic Journal: Microeconomics*, **6**, 35–58.

STROTZ, R. H. (1955), "Myopia and Inconsistency in Dynamic Utility Maximization", *The Review of Economic Studies*, **23**, 165.

TVERSKY, A. and KAHNEMAN, D. (1981), "The Framing of Decisions and the Psychology of Choice", *Science*, **211**, 453–458.

WEI, D. and GREEN, B. (2022), "(Reverse) Price Discrimination with Information Design" (SSRN Scholarly Paper ID 3263898, Social Science Research Network, Rochester, NY) 2020.

YU, P. C. (2020), "Seemingly Exploitative Contracts", *Journal of Economic Behavior & Organization*, **176**, 299–320.

——— (2022), "Optimal Retirement Policies with Present-Biased Agents", *Journal of the European Economic Association*, **19**, 2085–2130.

ZHANG, W. (2012), "Endogenous Preferences and Dynamic Contract Design", *The B.E. Journal of Theoretical Economics*, **12**, 19.