

forthcoming in *Review of Industrial Organization*

The Economics of Advice*

Winand Emons[†] Severin Lenhard[‡]

January 24, 2025

Abstract

A consumer wants to buy one of three different products. An expert observes which of the three products is the best match for the consumer. From his knowledge of costs and the observation of prices, the consumer can infer the expert's incentives. Under linear prices a monopolistic expert may truthfully reveal, may partially reveal, or may not reveal at all her information. The outcome is inefficient; moreover, the consumer gets some of the surplus. With a two-part tariff the expert truthfully reveals her information. The outcome is efficient, and the expert appropriates the entire surplus. If experts are competitive, they also truthfully reveal; here all of the surplus goes to consumers.

Keywords: advice, credence good, horizontal product differentiation.

JEL: D18, D82

*We thank Jean-Michel Benkert, Yuk-fai Fong, Byoung Heon Jun, Edi Karni, Rudi Kerschbamer, Jeong-Yoo Kim, Ting Liu, Oleg Muratov, Klaus Neusser, Alexander Rasch, Marco Schwarz, Markus Walzl, an anonymous referee, and Larry White for helpful comments. The usual disclaimer applies.

[†]Universität Bern and CEPR. Departement Volkswirtschaftslehre, Universität Bern, Postfach 8573, CH-3001 Bern, Switzerland, winand.emons@unibe.ch.

[‡]Universität St. Gallen and the Secretariat of the Swiss Competition Commission, severin.lenhard@unisg.ch.

1 Introduction

1.1 Motivation

When purchasing an experience good, consumers often do not know which variant of the product matches their preferences best. For instance, a consumer who wishes to buy a mattress may consider a foam, a spring, or a latex mattress. An investor may be interested in a short-term, an intermediate-term, or a long-term bond. A car buyer may consider a gasoline, a hybrid, or a diesel engine. Once the consumer has bought a variant, he will find out over time how well this product fits his needs. The consumer will, however, never discover whether one of the other two alternatives would have been a better match.

In such a situation consumers often rely on the advice of professional experts who know the products that they sell. The mattress dealer may, e.g., realize that the consumer who tends to sleep hot at night is best served with a spring mattress that offers better airflow than the alternatives. The investment advisor may find out that given the consumer's existing asset portfolio, intermediate bonds offer a good balance of portfolio diversification and expected returns over the long run. The car dealer may anticipate that the consumer values sporty rides, which makes a gasoline-powered engine the best fit.¹

Typically, the expert's advice constitutes a credence good (Darby and Karni (1973)): Even though the consumer observes ex post how well the recommended product fits his needs, he will never discover whether other alternatives would have served him better. To restate the definition of the different goods: The products that the consumer contemplates purchasing are experience goods.² The expert's recommendation about which of the

¹More examples abound: The gardener recommends ground cover plants, the sporting goods salesperson recommends running shoes, the plumber recommends a faucet, the travel agent recommends a hotel, the bookseller recommends a page-turner, ...

²These products can actually also be credence goods. For example, the intermediate-

experience goods best fits the consumer's needs is a credence good. Our focus is on the credence good *advice*.

To set the stage for the analysis of advice: The experts' information advantage creates an incentive to behave opportunistically. They may recommend products which are not in the best interest of consumers but in the best interest of the experts.³

Consequently, there is a need for mechanisms to generate truthful recommendations. Perhaps the simplest mechanism to ensure honest advice is the separation of recommendation and selling: The advising expert has no incentive to recommend unsuitable products, and the seller may only sell what has been recommended by her colleague. An example of this simple yet effective mechanism is the often encountered separation of the prescription and the sale of drugs: The medical doctor prescribes; the pharmacist sells. This "separation" mechanism fails, however, to do a good job when it is cheaper to provide advice and sale jointly rather than separately.

In this paper we examine a scenario where the expert gives advice and at the same time sells the products; visiting advice-only experts is simply too inconvenient. In such a scenario the "second-opinion" mechanism is also unattractive. Markets that we have in mind include consumer durables (cars, bicycles, electronics,...), financial assets, insurance, travel, etc. We analyze different market structures and pricing schemes. Linear prices are not good

term bond may have been indeed the best match for the investor from an ex ante perspective, although ex post the long-term bond performed better. We think, however, that the more common case is experience goods and stick to this interpretation.

³To give a few examples: There is much anecdotal evidence that the fee structure of investment products, rather than their suitability, drives their sale to customers; see Charter et al. (2010) and the references quoted therein. Anagol et al. (2017) show that life insurance agents in India commonly recommend products that do not cater to the consumers' needs, but rather increase commissions. The rate at which insurance agents recommend suitable products is found to be as low as 5%. The White House Council of Economic Advisers estimates that Americans lose about \$17 billion each year in foregone retirements earnings due to conflicted advice; see obamawhitehouse.archives.gov/the-press-office/2015/02/23/fact-sheet-middle-class-economics-strengthening-retirement-security-crac.

at generating truthful expert advice whereas two-part tariffs are.

1.2 Contribution

A consumer wants to buy one of three different products. At the time of purchase the consumer's future utility for each of the three products is uncertain. Utilities are equally distributed; and with equal probability each product can be the best match for the consumer.

The expert has more information than does the consumer. She observes which of the three products is the best fit for the consumer. Specifically, the expert observes which product will generate the highest utility.⁴ The knowledge of which product is the best fit increases the consumer's expected utility from buying this product; at the same time, it reduces his expected utility from buying one of the other two products.

The expert produces the three products at different costs. These different production costs generate the expert's vested interest: At the same price the seller has an incentive to sell the low-cost product to the consumer. The expert's costs are common knowledge: The knowledge of the costs and the observation of the prices allow the consumer to infer the expert's incentives.

We first examine the scenario where the expert is a monopolist that charges linear prices. The expert may follow any revelation strategy: She may lie; she may truthfully reveal all her information; or she may withhold information. In equilibrium she follows one of the following strategies: full revelation; partial revelation; and no revelation. Since partial revelation can arise only with a minimum of three products, we analyze three rather than two products.

In order to reveal truthfully all of her information, the expert has to be indifferent among the three products. The high-cost product yields the lowest mark-up. Therefore, she has to set the prices for the other two products such that they generate the same mark-up as the high-cost product. This implies

⁴She does, however, not observe how high the utility score is going to be.

that the consumer gets some of the surplus for the low-cost products.

If the expert reveals no information, consumers have the same willingness-to-pay for all products. Her profit is highest for the low-cost product. Accordingly, the expert charges prices that steer the consumer to the low-cost product.

Under partial revelation the expert reveals the medium-cost product and lumps the high- and the low-cost products together. When the consumer learns that he is of the high- or the low-cost type, his willingness-to-pay is higher than with no information. Prices are such that the consumer buys the low-cost version.

All three strategies may form an equilibrium. The higher is the cost of the expensive product, the more likely is the outcome to be inefficient: The expert lures the consumer away from the high-cost product by providing partial or no information. Moreover, even though the expert is a monopolist, some of the surplus goes to the consumer.

Next we allow the expert to charge two-part tariffs: In addition to the prices for the three products she charges a recommendation fee. She sets marginal cost prices for the products that make her indifferent as to her recommendation. She fully reveals and generates the entire surplus which, in turn, she appropriates with her recommendation fee.

This result immediately implies the outcome for competitive experts. They charge marginal cost prices. The recommendation fee is zero—unless there is a fixed cost of setting up shop; in this case the recommendation fee equals the fixed cost.

Linear prices are thus not an optimal tool when truthful expert advice is called for. For truthful reporting the high-cost product restricts the expert's price-setting range. Therefore, the expert lures the consumer away from the high-cost product by partial or no revelation. Moreover, we show that two-part tariffs are a powerful instrument to generate truthful advice: The product prices make the expert indifferent as to her recommendation, and

the recommendation fee serves to appropriate the surplus. This result holds for a monopoly and for competition.

1.3 Related Literature

There is a growing literature on credence goods and expert services; see Dulleck and Kerschbamer (2006), Balafoutas and Kerschbamer (2017), and Emons et al. (2024) for surveys. Most of this literature considers credence goods from a vertical product differentiation perspective. A consumer needs either a minor version of the good –typically a repair– that solves his problems in some contingencies; or he needs a major version of the good, which is always sufficient. This approach gives rise to the problems of under- or overtreatment: The expert performs only the minor repair although the major repair is necessary; or she performs the major treatment although the minor one would have been sufficient. We relate our findings to this literature in Section 4.⁵

A different strand of the credence goods literature examines financial advice from a horizontal product differentiation angle as we do; see Inderst and Ottaviani (2012) for a survey.⁶ Financial advisors possess more information than do their clients about which asset is the best fit for the consumer. Advisors act as brokers for upstream firms and are driven by financial incentives that are generated by commissions that are paid by product providers. Such a horizontally differentiated market gives rise to the problem of mistreatment/mis-selling: The advisor recommends a product that is not in the best interest of the consumer but in the best interest of the advisor.

⁵The literature also uses the terms under- or overprovision or supplier-induced demand. If the type of repair is not verifiable, the additional problem of overcharging arises: The expert performs the minor repair but charges for the major one.

⁶Consumers have different willingness-to-pay under horizontal and vertical differentiation. Under horizontal differentiation each consumer has his favorite product. Under vertical differentiation consumers with the minor problem value both types of repair equally because both repair types solve the minor problem. Consumers with the major problem value only the major repair.

The focus of this literature is on the role of the advisor as a firm's direct marketing agent: The upstream firm sets commissions to which the downstream advisor as the agent reacts. Our expert chooses prices/mark-ups herself. With her choice of prices she commits to a recommendation policy. Moreover, in Inderst and Ottaviani (2012) the financial advisor makes a dichotomous recommendation. By contrast, our expert may follow any revelation strategy: from lies, to the withholding of information, to the naked truth.

Bardey et al. (2020) consider a set-up with two horizontally differentiated products. In contrast to our framework, it is costly for the monopoly seller to collect information about the buyer's needs. The advisor invests only in this information if it affects her advice, which results in the truthful revelation of the entire information; in our set-up, the advisor may hold back some information.

Teh and Wright (2022) study a horizontal set-up with more than two products. Firms not only set commissions but also set product prices. The advisor ranks the products for the consumer. By paying an inspection cost, the consumer learns his match value from each product as well as its price. Teh and Wright (2022) thus consider search goods. The consumer knows the surplus before he buys. As in Inderst and Ottaviani (2012) the advisor is the marketing agent: Upstream firms set prices and commissions.⁷ More important: Since consumers neither observe prices nor commissions ex ante, these instruments cannot be used to commit to a recommendation policy as in our framework.

Finally, our paper is related to the literature on cheap talk and delegation; see, e.g., Crawford and Sobel (1982), Dessein (2002), or Chakraborty and Harbaugh (2010). Whereas in this literature the preference divergence of the sender and receiver is given exogenously, in our framework it arises

⁷In essence, the advisor is subject to resale-price maintenance, whereas our expert is free to set her prices.

endogenously through the expert's choice of prices.

The rest of this paper is organized as follows: The next section describes the model. In section 3 we derive the market outcomes. We first examine a monopolist that uses linear prices. Then we allow for two-part tariffs and consider competitive experts. In section 4 we discuss our findings. Section 5 concludes. Proofs and the derivation of the statistical results that are used throughout the paper are relegated to the Appendix.

2 Model

A consumer wants to buy one –and just one– among three different products: α , β , and γ . The consumer derives utility v_i from product $i \in \{\alpha, \beta, \gamma\}$. Since these are experience goods, the consumer learns the realization v_i only after consumption. At the time of purchase the consumer's utility for each of the three products is uncertain. Let the random variables V_i , $i \in \{\alpha, \beta, \gamma\}$ be independently and identically distributed. Moreover, with equal probability each of the three products can be the best match for the consumer. Not buying a product generates utility 0.

For ease of exposition –in particular to have numerical values for the various expectations with which we work– we assume the uniform distribution throughout the main text. At the end of this section we show that all of our results hold for any independently and identically distributed random variables. Accordingly, assume $V_i \sim U[0, 1]$, $i \in \{\alpha, \beta, \gamma\}$ and the three random variables are independent. The consumer thus expects $E(V_i) = 1/2$ from each product.

The products are sold by an expert. The expert has more information than the consumer. She observes which of the three products is the best fit for the consumer: The expert finds out that, say, product α will generate higher utility than will products β and γ . Formally, the expert observes $i = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}} \{v_j\}$. The expert, however, does not observe the realizations

of the utility scores: v_α may be higher than v_β and v_γ , yet v_α may still be pretty low. If good i is the best fit for the consumer, we call him to be of type $i \in \{\alpha, \beta, \gamma\}$.⁸

The knowledge of being type i increases the buyer's expected utility from purchasing the product: the probability that, say, v_α is higher than v_β and v_γ is fairly low for low values of v_α and fairly high for high values thereof. The distribution is thus left-skewed. Specifically, $E(V_i | i = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}} \{v_j\}) = 3/4$: The consumer's willingness-to-pay for good i increases by $1/4$ if he learns that he is of type $i \in \{\alpha, \beta, \gamma\}$.⁹ At the same time the expected utility of buying good $k \neq i$ decreases by $1/8$ when the consumer learns that he is of type i : $E(V_k | i = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}} \{v_j\}) = 3/8$.

If the expert partially reveals that the consumer is either of type i or j —that he is not type k —his expected utility of purchasing i or j increases by $1/16$: $E(V_i | k \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}} \{v_j\}) = 9/16$; at the same time his utility of buying good k decreases by $1/8$: $E(V_k | k \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}} \{v_j\}) = 3/8$.

The expert observes the consumer's type at no cost. Her cost of production are $c_\alpha := 0 < c_\beta < c_\gamma < 3/8$. These costs are common knowledge.¹⁰ All products yield positive surplus even if the consumer doesn't know his type. For the uninformed consumer, product α is the efficient choice: All products generate the same expected utility and α comes at the lowest cost. When informed, the consumer of type i should buy product i : $3/4 - c_i > 3/8 - c_j$ for $i, j, \in \{\alpha, \beta, \gamma\}$, $j \neq i$. In the efficient allocation the expert, therefore, informs the consumer about his type, and the consumer buys the corresponding product.

⁸In terms of our automobile example, the expert observes that the customer is, say, the sporty driver type, yet she cannot observe the extent to which he is going to use the vehicle. Note that our expert perfectly observes the consumer's type. Extending our set-up to an expert who observes the customer's type imperfectly yields qualitatively similar results.

⁹In Claims 1-5 in the Appendix we derive the different expected values that we use throughout the text for the uniform distribution.

¹⁰This assumption is crucial to our results. See the discussion in Section 4.

We have derived for the uniform distribution that

$$E(V_i|i = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) \geq E(V_i|k \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) \geq E(V_i) \geq E(V_i|i \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) = E(V_i|k = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}).$$

Claim 6 in the Appendix shows that this chain of weak inequalities holds for any independently and identically distributed variables V_α , V_β , and V_γ . The assumption $c_\gamma < 3/8$ —which ensures that the informed type i should buy product i in the uniform case—generalizes to $E(V_\gamma|\gamma = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) - E(V_\alpha|\gamma = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) > c_\gamma$. If this assumption on costs is satisfied, all of our results for the uniform scenario carry over to any independently and identically distributed variables.

3 Market Outcomes

In this section we examine market outcomes. We start with the scenario where the expert is a monopolist. Then we analyze competitive experts.

3.1 Monopolistic Expert

The three products are *experience goods*: The consumer learns the realization v_i only after consumption. The expert’s recommendation is a *credence good*: If the consumer buys product i , he observes ex post the utility score v_i , yet he never discovers whether product i was indeed his best match.¹¹

To recommend truthfully, the expert, therefore, has to be indifferent between the three products: The *equal compensation principle* has to hold.¹² The products yield, however, different surpluses, with product γ generating

¹¹The expert never observes the utility score v_i —neither ex ante nor ex post. A contractual arrangement—such as liability that is based on the realization v_i —is thus impossible. See Chen et al. (2022) for an analysis of liability in expert markets.

¹²We borrow this term from the principal-agent literature. See, e.g., Milgrom and Roberts (1992), 228-232. In the credence goods literature, Emons (1997, 2001, 2013) uses the term equal compensation principle, Dulleck and Kerschbamer (2006) use the term equal mark-up prices.

the smallest surplus $3/4 - c_\gamma$. Equal compensation implies that for truthful advice the expert can appropriate a maximum of $3/4 - c_\gamma$ for each product; thus, the consumer gets some of the surplus if he is of type α or β . In what follows we will show that the expert can do better by partially revealing her information or by revealing no information at all.

3.1.1 Linear Prices

Let us now describe the game: In the first stage the monopolistic expert sets prices p_i , $i \in \{\alpha, \beta, \gamma\}$. In the second stage the consumer visits the expert and observes prices. The expert observes the consumer's type and makes her recommendation. A recommendation is an element of the powerset without the empty set: $P(\{\alpha, \beta, \gamma\}) \setminus \emptyset$.¹³ Formally, the expert's recommendation strategy is $r : \{\alpha, \beta, \gamma\} \mapsto P(\{\alpha, \beta, \gamma\}) \setminus \emptyset$. The range of r forms the possible recommendations that a consumer receives, depending on his type. The expert maximizes expected profits. When she is indifferent as to her recommendation, she reports truthfully; she may, however, withhold information. In the event of indifference, she does not provide false information.¹⁴

The consumer maximizes his expected utility minus the price he pays. Upon receiving the recommendation, the consumer updates his beliefs about his type and eventually purchases a product or leaves the expert. If the recommendation is not profit-maximizing given the prices, the consumer infers a mistake and sticks to his priors. If there is no trade, both consumer and expert have a payoff of 0. We focus on perfect Bayesian equilibria.

In Lemma 1 we derive the profit maximizing prices that support truthful revelation of the consumer's type. Prices are such that all three products have the same mark-up. With these prices the monopolist does not appropriate the entire surplus: The most costly product γ determines her mark-up.

¹³We assume that the messages \emptyset or $\{\alpha, \beta, \gamma\}$ amount to the same thing: The consumer gets no new information and we drop \emptyset from the message space. Alternatively, because $\operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}} \{v_j\}$ always exists in our setup, \emptyset cannot be a truthful message.

¹⁴Upon observing, say, α , the expert may thus report $\{\alpha\}$, $\{\alpha, \beta\}$, $\{\alpha, \gamma\}$, $\{\alpha, \beta, \gamma\}$.

Lemma 1: *Suppose $p_\alpha = 3/4 - c_\gamma$, $p_\beta = 3/4 - (c_\gamma - c_\beta)$, $p_\gamma = 3/4$ and let the consumer be of type $i \in \{\alpha, \beta, \gamma\}$. The expert truthfully reveals i : $r(i) = \{i\}$. The consumer believes the message and purchases product i . These prices are profit maximizing within the set of fully revealing prices.*

In Lemma 2 we derive prices under which the expert reveals no information: With no new information, the consumer sticks to his priors. Since he expects the same surplus from all three products, he goes for the cheapest one. If the consumer has the same willingness-to-pay for all three products, the expert steers the consumer away from β and γ and charges $1/2$ for α . Accordingly, here it is the least expensive product α that determines her mark-up.

Lemma 2: *Suppose $p_\alpha = 1/2$, $p_\beta > 1/2$, $p_\gamma > 1/2$ and let the consumer be of type $i \in \{\alpha, \beta, \gamma\}$. The expert reports $\{\alpha, \beta, \gamma\}$: $r(i) = \{\alpha, \beta, \gamma\}$. The consumer sticks to his priors and buys product α . These prices are profit maximizing within the set of no revealing prices.*

In Lemma 3 we derive prices where the expert partially reveals the type: She pools the highest-cost product with the least-cost product. The mark-ups for the reports $\{\alpha\}$, $\{\beta\}$, and $\{\alpha, \gamma\}$ are equal and higher than for $\{\gamma\}$: When the consumer hears $\{\alpha, \gamma\}$, he expects $9/16$ and buys α ; when he hears $\{\beta\}$, he expects $3/4$ and buys β .

Why does the expert pool product γ with product α and not with product β ? Type γ 's willingness-to-pay for product α following the recommendation $\{\alpha, \gamma\}$ is the same as his willingness-to-pay for product β after the recommendation $\{\beta, \gamma\}$. The expert has lower costs if type γ buys product α rather than β . Therefore, she steers the most expensive type to the least expensive product.

Lemma 3: *Suppose $p_\alpha = \min\{9/16, 3/4 - c_\beta\}$, $p_\beta = \min\{9/16 + c_\beta, 3/4\}$, and $p_\gamma \in (p_\alpha, p_\alpha + c_\gamma]$; and let the consumer be of type $i \in \{\alpha, \beta, \gamma\}$. The*

expert reports $\{\alpha, \gamma\}$ when the consumer is of type α or γ ; she reports $\{\beta\}$ when the consumer is of type β . Formally, $r(s) = \{\alpha, \gamma\}$; $s \in \{\alpha, \gamma\}$; and $r(\beta) = \{\beta\}$. The consumer believes the reports and buys product α when he hears $\{\alpha, \gamma\}$ and buys product β when he hears $\{\beta\}$. These prices are profit maximizing within the set of partially revealing prices.

We may now establish equilibria for the entire game.

Proposition 1:

- (i) For $c_\gamma \leq 3/16$, the monopolist sets prices and truthfully reveals all information as is described in Lemma 1;
- (ii) for $c_\beta \geq 1/4$, the monopolist sets prices and reveals no information as is described in Lemma 2;
- (iii) for $c_\gamma \geq 3/16$ and $c_\beta \leq 1/4$, the monopolist sets prices and partially reveals information as is described in Lemma 3.

The proof follows from simply comparing the expert's profits in the three constellations. Figure 1 shows which products are traded in equilibrium. In the vertically shaded area the monopolist truthfully reveals all information. All products are traded, and the resulting allocation is efficient. In the diagonally shaded area the monopolist partially reveals her information. The consumer never buys products γ , which is inefficient when he is of this type. In the horizontally shaded area the expert reveals no information at all. Only product α is traded, which is inefficient when the consumer is of type β or γ . Overall, unless c_γ is small, the equilibrium is, therefore, inefficient.

The higher is the consumer's valuation, the more that the expert can charge. The consumer's willingness-to-pay is highest when the expert reveals truthfully. To reveal truthfully, the expert has to be indifferent between recommending any of the three products: The equal compensation principle

has to hold. The maximal mark-up the expert can charge is determined by the most expensive product, which generates a profit of $3/4 - c_\gamma$.

Increasing c_γ makes full revelation less attractive. The expert wants to lure the consumer away from good γ . To do so, she sets the price p_γ such that the product is not attractive for the consumer and her mark-up is small; this reduces her financial incentives to report $\{\gamma\}$.

Under partial revelation she pools product γ with product α ; the consumer buys the cheaper of the two products $-\alpha-$ which generates the profit $9/16$ or $3/4 - c_\beta$, whichever is lower. She reveals β , which generates the profit $\min\{9/16, 3/4 - c_\beta\}$.

insert **Figure 1** around here

Partial revelation becomes unattractive with high c_β . The expert lures the consumer away from β and γ by pooling all three products. With no additional information, the consumer purchases the cheapest product $-\alpha-$ which generates the profit $1/2$.

3.1.2 Two-part Tariffs

The equilibrium allocation with linear prices is inefficient. Furthermore, the expert does not appropriate the entire surplus that is generated: Under full revelation the consumer gets some surplus when he is of type α or β , and also under partial revelation when he is of type β . Only under no revelation is the expected surplus of the consumer zero.¹⁵

Both phenomena are due to the fact that the expert is restricted to use linear prices. Once we allow for two-part tariffs, efficiency is attained, and the monopolist appropriates the entire surplus.

¹⁵Accordingly, among the equilibrium allocations the consumer prefers full to partial to no revelation.

Consider the following modification of our game: In the first stage the expert sets p_r –the price that she charges for her recommendation– together with p_i , $i \in \{\alpha, \beta, \gamma\}$: the prices for the three products. In the second stage the consumer observes the prices and decides whether he wants the expert’s recommendation for the price p_r . If the consumer visits the expert, the game continues as our previous one:

Proposition 2: *Let $p_r = 3/4 - (p_\alpha + p_\beta + p_\gamma)/3$, $p_\alpha = p_\gamma - c_\gamma$, $p_\beta = p_\gamma - (c_\gamma - c_\beta)$, $p_\gamma \in [c_\gamma, 3/4]$, and let the consumer be of type $i \in \{\alpha, \beta, \gamma\}$. The consumer consults the expert and pays p_r . The expert truthfully reveals i : $r(i) = \{i\}$. The consumer believes the message and purchases product i .*

With this two-part tariff, product prices take care of equal compensation. The expert implements the efficient allocation and appropriates the entire surplus that is thus generated with the recommendation fee. Note that $p_\gamma = 3/4$ results in the product prices of Lemma 1; and $p_\gamma = c_\gamma$ generates marginal cost prices.

3.2 Competitive Experts

For a set-up with competitive experts Proposition 2 immediately implies the following Corollary. Let there be $l = 1, \dots, L$, $L \geq 2$, experts who simultaneously charge prices p_r^l , p_i^l , $i \in \{\alpha, \beta, \gamma\}$ in stage 1. In the second stage the consumer observes the list of prices and then decides whether or not to visit an expert. If he does so, the game continues as our previous ones:

Corollary: *Let $p_r^l = 0$, $p_\alpha = 0$, $p_\beta = c_\beta$, $p_\gamma = c_\gamma$ for $l = 1, \dots, L$ and let the consumer be of type $i \in \{\alpha, \beta, \gamma\}$. The consumer consults an expert. The expert truthfully reveals i : $r(i) = \{i\}$. The consumer believes the message and purchases product i .*

Marginal cost pricing ensures equal compensation, and Bertrand competition drives profits (or p_r) down to zero. For this result to hold we do not

need a two-part tariff: Linear prices also do the job.

Our last two results show that marginal-cost prices solve the experts' credibility problem. In a competitive framework where profits are zero, linear marginal-cost prices provide proper incentives to recommend the correct product. If the expert has market power as in our framework, or when she has a fixed cost that she needs to recover, we need in addition the recommendation fee.

4 Discussion

Our results immediately beg a few questions. First, under linear prices the high-cost products β and γ restrict the expert's price-setting range. Why doesn't the expert stop offering these products? Specifically, in the no-revelation equilibrium the expert steers the consumer away from products β and γ by charging high prices; she could as well drop these products from her product line. Likewise, in the partial revelation equilibrium the expert could stop offering γ .

In our set-up it makes no difference whether the expert has a broad range of products and doesn't sell some of these or whether she offers a narrow product range to begin with: There is no fixed cost that is associated with offering a product. By contrast, with such a fixed cost under partial revelation the expert would offer only α and β ; and under no revelation she would restrict herself to selling α .

Our results are thus compatible with experts who offer the full product range –yet charging sky-high prices for some of their products– and experts who offer a narrow product range. For example, under no revelation our expert could offer only one product like a discounter. Nonetheless, whereas the typical discounter has a limited product range so as to reduce costs, our expert drops certain items in order to avoid cannibalizing the margins of the products that she plans to sell.

Offering the full product range may make sense for reasons beyond our framework. For example, the marketing literature asserts that a deep rather than a shallow assortment may give consumers the impression of having a choice, which increases the likelihood that they shop at the store in the first place; see, e.g., Broniarczyk et al. (1998).¹⁶

Second, why do we not observe more use of two-part tariffs when expert advice is called for? Although there is wide consensus that kickbacks lead to mis-selling, volume-based sales commissions remain the most widespread form of remuneration for advice in insurance, credit, and investment markets.

Finland, Denmark, the UK, and the Netherlands introduced a ban on commissions so that brokers had to switch to recommendation fees. The experience with (hourly) recommendation fees has been at best mixed. In Finland consumers were unwilling to pay the high fees of intermediaries. In the UK customers had to pay more for financial advice than they did before the ban. In addition, following the ban, the number of brokers or advisers declined in Finland, Denmark, and the UK (Reifner et al. (2013)).

Several explanations may reconcile our results with these empirical observations: A separate recommendation fee increases transparency: Consumers become aware of how expensive advice actually is and may decide that it is not worth it. Furthermore, with a separate recommendation fee consumers can reject the recommendation, yet still buy the product relying solely on their priors. If the recommendation charge is included in the product price, they don't have this option: They can only buy the entire package.

In our framework consumers know that they will buy one of the three products after observing prices: The recommendation always results in a successful purchase. However, suppose that there is the possibility that consumers may not want to buy after receiving the recommendation. Then the payment of the recommendation fee results in a negative payoff. In such a

¹⁶This is one explanation as to why discounters add national brands to their private labels.

set-up a loss-averse consumer may turn down the opportunity to obtain the recommendation that an expected utility maximizer would buy.¹⁷

Two-part tariffs weakly dominate linear prices in our framework. Yet, when the costs of the three products are similar –if c_γ is small– the expert truthfully reveals all of her information under linear prices.¹⁸ The outcome is efficient. Moreover, when c_γ decreases, total surplus increases and so does the share thereof that accrues to the expert. Therefore, if c_γ is small, the expert gains little by switching from linear prices to a two-part tariff: Her incentive to use a two-part tariff is lower when the products’ costs are closer to each other. Furthermore, competitive experts may use either linear prices or two-part tariffs: They make zero profits in either case. Accordingly, for a richer set-up our results suggest that linear prices are more likely to be observed if costs are similar or if the market is competitive.

Our consumers are identical in that they obtain the same expected surplus. The expert appropriates the entire surplus through the recommendation fee p_r . If there are, say, two groups of consumers who differ with respect to the expected surplus, the expert can no longer appropriate the entire surplus with our simple two-part tariff. If p_r equals the expected surplus of the high-surplus group, the expert loses the low-surplus group. If p_r equals the expected surplus of the low-surplus group, the expert serves both groups, yet does not capture the entire surplus of the high-surplus group. Analyzing two-part tariffs and linear prices with heterogeneous consumers is perhaps an interesting topic for future research.¹⁹

Note that our two-part tariff scenario corresponds to the commitment case in the vertical differentiation set-up of Dulleck and Kerschbamer (2006):

¹⁷In an online experiment Chater et al. (2010) find people disproportionately averse to paying an up-front fee for advice. Between 20-30% of the subjects displayed evidence of narrow framing and loss aversion, which is consistent with their being excessively averse to an up-front fee.

¹⁸Recall that $c_\alpha := 0 < c_\beta < c_\gamma$.

¹⁹If the recommendation fee has to be uniform across the heterogeneous consumers, we are in the world of Oi’s (1971) “Disneyland Dilemma”.

There is no recommendation fee; however, once the expert makes a recommendation, the customer is committed to undergo the recommended treatment—even if his surplus thereof is negative. Under our two-part tariff type γ ex post ends up with negative surplus and regrets to have consulted the expert in the first place. With commitment in Dulleck and Kerschbamer (2006) the outcome is efficient, and the expert captures the entire surplus, as is the case under our two-part tariff.

Fong et al. (2014) study the no-commitment set-up: There is no recommendation fee. The client can reject treatment recommendations. The expert has to set prices that yield a non-negative client surplus for each treatment. The expert trades off efficiency versus rent extraction. No-commitment corresponds to our linear prices framework; it also results in inefficiencies and less than full surplus appropriation by the expert. Our results, therefore, corroborate findings of the vertical credence goods literature.²⁰

Third, how is our two-part tariff related to the “separation” mechanism that we mentioned in the Introduction? For truthful revelation under our two-part tariff, product prices generate equal mark-ups that render the expert indifferent as to her recommendation. With the recommendation fee she appropriates the surplus. The separation mechanism does not need such fine-tuning of prices. Any product prices above marginal costs do the job. Even if the seller has a strong incentive to sell, say, product β , she cannot do so unless β has been suggested by her advising colleague who, in turn, has no monetary incentive to recommend a particular product.

Fourth, what is the outcome with more than three products? We have (in unpublished work) analyzed the case of four products α, β, γ , and δ —with product δ having the highest cost. The results are qualitatively similar. If c_δ is small, the expert fully reveals. If c_β is large, she does not reveal at all. Otherwise, she partially reveals. Partial revelation is clearly more complicated with four products: The number of possible messages more than

²⁰Hu and Lei (2023) perform a similar exercise as we do in a vertical framework.

doubles for each product that we add to the assortment.²¹ We conjecture that this general structure also holds for more than four products.

Fifth, the assumption that costs are common knowledge is crucial for our results. This assumption seems reasonable in quite a few contexts. For example, disclosure of financial interests may be mandatory. The Physician Payments Sunshine Act (PPSA) requires physicians to disclose financial interests in manufacturers of drugs, devices, biologicals, and medical supplies. Similarly, the U.S. Securities and Exchange Commission (SEC) issues guidance that requires that investment advisors disclose financial conflicts of interest when advising clients. So does Art. 28, Chapter VI, Directive (EU) 2016/97, which requires that insurance intermediaries and insurance undertakings take all appropriate steps to identify conflicts of interest between themselves and their customers.

The assumption that consumers observe the expert's costs is the standard one in the credence goods literature.²² It allows us to compare our findings to the results of this literature. In principle, the expert has an incentive to reveal her costs: Knowing the costs, the consumer can compute the mark-ups and, therefore, infer the expert's incentives. This, in turn, allows the expert to meaningfully use recommendations to maximize her profits. To put this differently: If consumers do not observe the costs, the expert tends to be stuck with the no revelation equilibrium and forgoes the profit opportunities of partial and full revelation.

5 Conclusions

In this paper we study experts who advise consumers as to which product to buy. Under linear prices a monopolistic expert may truthfully reveal, may partially reveal, and may not reveal at all her information. The outcome is inefficient; moreover, the consumer gets some of the surplus.

²¹Since we neglect \emptyset , the number of possible messages with N products is $2^N - 1$.

²²An exception is Hilger (2016).

Efficiency is restored with a two-part tariff: In addition to product prices the expert charges a recommendation fee. The product prices render the expert indifferent as to her recommendation; the recommendation fee serves to appropriate the surplus. If experts are competitive, they also truthfully reveal under a two-part tariff; here all the surplus goes to consumers. Two-part tariffs are, therefore, a clever instrument to generate truthful advice.

When we discuss the expert, we have a real person in mind: We use the pronoun *she* after all: *She* knows her products; *she* interacts with the consumer; and eventually *she* finds out the consumer's best match.

Nevertheless, our framework applies also to online recommendation systems. The seller can enter product information as well as all of the data that she has collected about the customer into the system. A clever algorithm combines these two data sets and develops an individual recommendation for the consumer. Picking up on our bond example: If the algorithm knows, e.g., the consumer's age, it seems sensible not to recommend a long-term bond to a ninety-year-old person. Yet, the algorithm can as well ignore the plethora of available data and simply recommend the product with the highest profit margin.

Appendix

In this Appendix we first prove the results that are provided in the text. Then we derive the expected values that we use throughout the text. Finally, we show that the ranking of the expected values holds for any independently and identically distributed variables.

Proof of Lemma 1: To reveal fully the consumer's type, i.e., $r(i) = \{i\}$, the expert has to be indifferent between the three reports $\{\alpha\}$, $\{\beta\}$, and $\{\gamma\}$: We must have $p_\alpha = p_\beta - c_\beta = p_\gamma - c_\gamma$. With these prices the expert has no incentive to report $j \neq i$. The consumer believes the message and has willingness-to-pay $3/4$ for product i . Setting $p_\gamma = 3/4$ yields the profit-maximizing prices with this property.

Consider now deviations to the partially revealing strategies $r(i) = \{i, j\}$, for $i, j \in \{\alpha, \beta, \gamma\}, j \neq i$, and to the unrevealing strategy $r(i) = \{\alpha, \beta, \gamma\}$. Whatever

the consumer believes when he hears one of these messages, his willingness-to-pay cannot exceed $3/4$. Accordingly, the expert cannot do better by deviating to one of those reports. \square

Proof of Lemma 2: With the message $\{\alpha, \beta, \gamma\}$ the consumer gets no new information and sticks to his priors. Since he expects $1/2$ from all three products with equal probability, he chooses the cheapest one α . If the consumer has a willingness-to-pay of $1/2$ for all three products, it is optimal for the expert to steer the consumer away from β and γ and to charge $1/2$ for α .

Consider now deviations to strategies with a single recommendation range. The consumer, thus, receives the message $\{\alpha\}$, $\{\beta\}$, or $\{\gamma\}$. Suppose the mark-ups of the three products are all different. Then the expert always recommends the product with the highest mark-up: say, i . Given the strategy of always recommending i , the consumer sticks to his priors when he hears this report. If the consumer hears $j \neq i$, he infers a mistake and also sticks to his priors. The consumer buys α , which makes a deviation not attractive.

If all mark-ups equal $1/2$, the expert reveals the truth. The consumer buys the recommended product, which generates a mark-up of (by assumption) $1/2$. A deviation is thus not attractive.

Suppose the mark-ups for β and γ are equal and greater than $1/2$. The expert is indifferent between recommending β and γ and prefers either report to α . When the consumer hears $\{\beta\}$ ($\{\gamma\}$), he knows that he is of type α or β (α or γ) and expects $9/16$. He buys α because it is the cheaper one of the two products. This deviation is thus not attractive.

If the mark-up for β (γ) equals $1/2$ while the mark-up for γ (β) exceeds $1/2$, the expert always recommends γ (β). The consumer sticks to his priors and accordingly buys α .

Consider now deviations to strategies with a partially revealing recommendation: $\{\alpha, \beta\}$, $\{\alpha, \gamma\}$, $\{\beta, \gamma\}$. Suppose that the mark-ups of the three products are all different. Then the expert always recommends the product with the highest mark-up: say, i . A partially revealing strategy is thus a mistake. The consumer sticks to his priors and buys α .

If all mark-ups equal $1/2$, the partially revealing message is truthful. The consumer expects $9/16$ and buys the cheaper of the two products. The mark-up is $1/2$, which makes this deviation not attractive.

Suppose the mark-ups for β and γ are equal and greater than $1/2$. When the consumer hears $\{\alpha, \beta\}$, $\{\alpha, \gamma\}$, he knows he is of type α or β (α or γ) and expects $9/16$. He buys α because it is the cheaper of the two products. When the consumer hears $\{\beta, \gamma\}$, he sticks to his priors and buys α . This deviation is thus not attractive. \square

Proof of Lemma 3: When the consumer hears $\{\alpha, \gamma\}$, he expects $9/16$ and buys α ; when he hears $\{\beta\}$, he expects $3/4$ and buys β . The mark-ups that accompany the reports $\{\alpha\}$, $\{\beta\}$, and $\{\alpha, \gamma\}$ are equal and higher than for $\{\gamma\}$.

These prices are profit maximizing within the set of partially revealing prices. Suppose the that expert pools products $\{i, j\}$ and truthfully reveals type $\{k\}$. The willingness-to-pay of type k equals $3/4$; the willingness-to-pay of the other two types equals $9/16$. Let i be the product that the expert intends to sell to the pooled types. Hence, the mark-up of k and i have to be equal so that the equal compensation principle holds.

On the one hand, the expert can extract the entire surplus of type k , which results in a mark-up of $3/4 - c_k$; on the other hand, the expert can extract the entire surplus of type i , which results in a mark-up of $9/16 - c_i$. However, the equal compensation principle implies that the expert can realize only the minimum of the two mark-ups: $\min\{3/4 - c_k, 9/16 - c_i\}$ weakly decreases in c_i and c_k . This implies that $j = \gamma$: The expert never sells the most expensive product. Next, note $\min\{3/4 - c_\beta, 9/16\} \geq \min\{3/4, 9/16 - c_\beta\}$. Thus, $i = \alpha$ and $k = \beta$ maximizes the expert's profit.

Let $p_\gamma < p_\alpha + c_\gamma$. Consider now deviations to strategies with a single recommendation range. The consumer, thus, gets the message $\{\alpha\}$, $\{\beta\}$, or $\{\gamma\}$. When the consumer hears $\{\alpha\}$, ($\{\beta\}$), he knows he is of type $\{\alpha, \gamma\}$, ($\{\beta, \gamma\}$) and has a willingness-to-pay of $9/16$. In the first case he buys the cheaper product α ; in the second case he does not buy at all. If he hears $\{\gamma\}$, he infers a mistake, sticks to his priors, and does not buy. These deviations are thus not profitable.

If the expert deviates to $r(\alpha) = \{\alpha\}$ and $r(s) = \{\beta, \gamma\}$ ($s \in \{\beta, \gamma\}$), the consumer expects $9/16$ and does not buy. If the expert deviates to $r(s) = \{\alpha, \beta\}$ ($s \in \{\alpha, \beta\}$), and $r(\gamma) = \{\gamma\}$, the consumer expects $9/16$ and buys α . These partially revealing strategies are thus not profitable.

Consider the deviation to the unrevealing strategy $r(i) = \{\alpha, \beta, \gamma\}$. With the message $\{\alpha, \beta, \gamma\}$ the consumer gets no new information and sticks to his priors. He expects $1/2$ from all three products and does not buy.

Finally, if $p_\gamma = p_\alpha + c_\gamma$, the expert makes the same mark-up with all three products; thus, no profitable deviation exists. \square

Proof of Proposition 2: The mark-ups of all three products are equal. Given that the consumer visits the expert, she recommends truthfully; and the consumer purchases the recommended product. Ex ante the consumer expects a surplus of $3/4 - (p_\alpha + p_\beta + p_\gamma)/3$, which is equal to the price that the expert charges for her recommendation. \square

Derivations of the expected values for the uniform distribution

Let $V_i \sim U[0, 1]$, $i \in \{\alpha, \beta, \gamma\}$. Moreover, the three random variables are independent.

Claim 1: $E(V_i) = 1/2$.

Proof: By definition $E(V_i) = \int_0^1 v \, dv = 1/2$. \square

Claim 2: $E(V_i | i = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}} \{v_j\}) = 3/4$.

Proof: Note that $E(V_i | i = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}} \{v_j\}) = E(\max\{V_i, V_j, V_k\})$. Thus, we derive the distribution of $\max\{V_i, V_j, V_k\}$. Independence implies $P(\max_i \{V_i\} < v) = \prod_i P(V_i < v)$. Moreover, because the distributions are identical, $F(v) = P(V_i < v)$, which results in the cumulative distribution function $F(v)^N$ for N identically and independently distributed variables. Accordingly, the probability density function is $NF(v)^{N-1}f(v)$, and

$$\begin{aligned} E(\max_i \{V_i\}) &= \int_0^1 v N F(v)^{N-1} f(v) dv \\ &= [v F(v)^N]_0^1 - \int_0^1 F(v)^N dv \\ &= 1 - \int_0^1 F(v)^N dv. \end{aligned}$$

For the uniform distribution we obtain the closed-form solution $1 - 1/(N + 1) = N/(N + 1)$ and for $N = 3$ this yields $3/4$. \square

Claim 3: $E(V_i | i \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}} \{v_j\}) = 3/8$.

Proof: By the law of iterated expectation

$$\begin{aligned} E\left(E(V_i | i \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}} \{v_j\})\right) &= E(V_i) \\ \frac{1}{3}E(V_i | i = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}} \{v_j\}) + \frac{2}{3}E(V_i | i \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}} \{v_j\}) &= E(V_i) \\ \frac{1}{3} \frac{3}{4} + \frac{2}{3}E(V_i | i \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}} \{v_j\}) &= \frac{1}{2} \\ E(V_i | i \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}} \{v_j\}) &= \frac{3}{8}. \quad \square \end{aligned}$$

Claim 4: $E(V_i | k \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}} \{v_j\}) = 9/16$.

Proof: Note that the probability density function can be written as $(\mathbb{1}(v_i \geq v_k) + \mathbb{1}(v_j \geq v_k) - \mathbb{1}(v_i \geq v_k)\mathbb{1}(v_j \geq v_k))/(2/3)$, where $\mathbb{1}$ is the indicator function. Thus,

$$E(V_i|k \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) = \frac{3}{2} \left(\int_0^1 \int_{v_k}^1 v_i \, dv_i \, dv_k + \int_0^1 \int_{v_k}^1 \int_0^1 v_i \, dv_i \, dv_j \, dv_k - \int_0^1 \int_{v_k}^1 \int_{v_k}^1 v_i \, dv_i \, dv_j \, dv_k \right) = 9/16. \quad \square$$

Claim 5: $E(V_i|k = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) = 3/8$.

Proof: By the law of iterated expectation

$$\begin{aligned} E\left(E(V_i|k = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\})\right) &= E(V_i) \\ \frac{2}{3}E(V_i|k \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) + \frac{1}{3}E(V_i|k = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) &= E(V_i) \\ \frac{2}{3} \frac{9}{16} + \frac{1}{3}E(V_i|k = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) &= \frac{1}{2} \\ E(V_i|k = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) &= \frac{3}{8}. \quad \square \end{aligned}$$

Ranking of the expected values for any i.i.d. variables

Claim 6:

$$\begin{aligned} E(V_i|i = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) &\geq E(V_i|k \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) \geq E(V_i) \geq \\ E(V_i|i \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) &= E(V_i|k = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}). \end{aligned}$$

for any independently and identically distributed variables V_α, V_β , and V_γ .

Proof: For any three random variables, there exist $3! = 3*2*1 = 6$ different cases to order the outcomes. Let us define the subset of outcomes for each different case; let $o_1 = \{v_i, v_j, v_k | v_i \geq v_j \geq v_k\}$, $o_2 = \{v_i, v_j, v_k | v_i \geq v_k \geq v_j\}$, $o_3 = \{v_i, v_j, v_k | v_j \geq v_i \geq v_k\}$, $o_4 = \{v_i, v_j, v_k | v_j \geq v_k \geq v_i\}$, $o_5 = \{v_i, v_j, v_k | v_k \geq v_i \geq v_j\}$, and $o_6 = \{v_i, v_j, v_k | v_k \geq v_j \geq v_i\}$.

Note that $i = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}$ implies the ordering in o_1 or o_2 , which are both equally likely due to the independently and identically distributed variables. Similarly, $k \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}$ implies the ordering in o_1, o_2, o_3 , or o_4 , of which are again all equally likely.

With this, we can rewrite the expectations, where we use $\mathbf{V} = [v_i, v_j, v_k] \subseteq \mathbb{R}^3$ as the vector of our three random variables:

$$\begin{aligned}
E(V_i|i = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) &= \\
&\frac{1}{2}E(V_i|\mathbf{V} \in o_1) + \frac{1}{2}E(V_i|\mathbf{V} \in o_2), \\
E(V_i|k \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) &= \\
&\frac{1}{4}E(V_i|\mathbf{V} \in o_1) + \frac{1}{4}E(V_i|\mathbf{V} \in o_2) \\
&+ \frac{1}{4}E(V_i|\mathbf{V} \in o_3) + \frac{1}{4}E(V_i|\mathbf{V} \in o_4), \\
E(V_i) &= \frac{1}{6}E(V_i|\mathbf{V} \in o_1) + \frac{1}{6}E(V_i|\mathbf{V} \in o_2) \\
&+ \frac{1}{6}E(V_i|\mathbf{V} \in o_3) + \frac{1}{6}E(V_i|\mathbf{V} \in o_4) \\
&+ \frac{1}{6}E(V_i|\mathbf{V} \in o_5) + \frac{1}{6}E(V_i|\mathbf{V} \in o_6), \\
E(V_i|i \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) &= \\
&\frac{1}{4}E(V_i|\mathbf{V} \in o_3) + \frac{1}{4}E(V_i|\mathbf{V} \in o_4) \\
&+ \frac{1}{4}E(V_i|\mathbf{V} \in o_5) + \frac{1}{4}E(V_i|\mathbf{V} \in o_6), \\
E(V_i|k = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) &= \\
&\frac{1}{2}E(V_i|\mathbf{V} \in o_5) + \frac{1}{2}E(V_i|\mathbf{V} \in o_6).
\end{aligned}$$

For the expectations of $E(V_i|\cdot)$, j respectively k are labels that can be interchanged because the variables are independently and identically distributed. Thus, conditional on o_1 yields the same result as conditional on o_2 . Similarly, conditional on o_3 yields the same as conditional on o_5 , and conditional on o_4 yields the same as conditional on o_6 : Only the relative ranking matters:

$$\begin{aligned}
E(V_i|i = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) &= E(V_i|\mathbf{V} \in o_1), \\
E(V_i|k \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) &= \frac{1}{2}E(V_i|\mathbf{V} \in o_1) \\
&+ \frac{1}{4}E(V_i|\mathbf{V} \in o_3) + \frac{1}{4}E(V_i|\mathbf{V} \in o_5), \\
E(V_i) &= \frac{1}{3}E(V_i|\mathbf{V} \in o_1) + \frac{1}{3}E(V_i|\mathbf{V} \in o_3) \\
&+ \frac{1}{3}E(V_i|\mathbf{V} \in o_5), \\
E(V_i|i \neq \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) &= \\
&\frac{1}{2}E(V_i|\mathbf{V} \in o_3) + \frac{1}{2}E(V_i|\mathbf{V} \in o_5), \\
E(V_i|k = \operatorname{argmax}_{j \in \{\alpha, \beta, \gamma\}}\{v_j\}) &= \\
&\frac{1}{2}E(V_i|\mathbf{V} \in o_3) + \frac{1}{2}E(V_i|\mathbf{V} \in o_5).
\end{aligned}$$

Since $E(V_i|\mathbf{V} \in o_1) \geq E(V_i|\mathbf{V} \in o_3) \geq E(V_i|\mathbf{V} \in o_5)$, and all of the expressions above are a weighted average of these three terms, the claim immediately follows. \square

References

- ANAGOL, SANTOSH, SHAWN COLE, AND SHAYAK SARKAR. 2017. “Understanding the advice of commissions-motivated agents: Evidence from the Indian life insurance market,” 99 *Review of Economics and Statistics*, 1-15.
- BALAFOUTAS, LOUKAS AND RUDOLF KERSCHBAMER. 2020. “Credence goods in the literature: what the past fifteen years have taught us about fraud, incentives, and the role of institutions,” 26 *Journal of Behavioral and Experimental Finance*, 100285
- BARDEY, DAVID, DENIS GROMB, DAVID MARTIMORT, AND JÉRÔME POUYET. 2020. “Controlling Sellers who Provide Advice: Regulation and Competition,” 68 *Journal of Industrial Economics*, 409-444.
- BRONIARCZYK, SUSAN, WAYNE HOYER, AND LEIGH MCALISTER. 1998. “Consumers’ perceptions of the assortment offered in a grocery category: The impact of item reduction,” 35 *Journal of Marketing Research*, 166-176.
- CHAKRABORTY, ARCHISHMAN AND RICK HARBAUGH. 2010. “Persuasion by Cheap Talk,” 100 *American Economic Review*, 2361-2382.
- CHATER, NICK, STEFFEN HUCK, AND ROMAN INDERST. 2010. “Consumer Decision-Making in Retail Investment Services: A Behavioural Economics Perspective,” Report to the European Commission/SANCO.
- CHEN, YONGMIN, JIANPEI LI, AND JIN ZHANG. 2022. “Efficient Liability in Expert Markets,” 63 *International Economic Review*, 1717-1744.
- CRAWFORD, VINCENT AND JOEL SOBEL. 1982. “Strategic Information Transmission,” 50 *Econometrica*, 1431-51.
- DARBY, MICHAEL AND EDI KARNI. 1973. “Free Competition and the Optimal Amount of Fraud,” 16 *Journal of Law and Economics*, 67-88.
- DESSEIN, WOUTER. 2002. “Authority and Communication in Organizations,” 69 *Review of Economic Studies*, 811-38.
- DULLECK, UWE AND RUDOLF KERSCHBAMER. 2006. “On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods,” 44 *Journal of Economic Literature* 5-42.
- EMONS, WINAND. 1997. “Credence Goods and Fraudulent Experts,” 28 *Rand Journal of Economics*, 107-119.
- EMONS, WINAND. 2001. “Credence Goods Monopolists,” 19 *International Journal of Industrial Organization*, 375-389.
- EMONS, WINAND. 2013. “Incentive Compatible Reimbursement Schemes for Physicians,” 169 *Journal of Institutional and Theoretical Economics*, 605-620.

- EMONS, WINAND, RUDOLF KERSCHBAMER, AND GERD MÜHLHEUSSER. 2024. “Credece Goods: A Symposium,” 180 *Journal of Institutional and Theoretical Economics*, 1-16.
- FONG, YUK-FAI, TING LIU, AND DONALD WRIGHT. 2014. “On the role of verifiability and commitment in credence goods markets,” 37 *International Journal of Industrial Organization*, 118-129.
- HILGER, NATHANIEL. 2016. “Why don’t consumers trust experts? Credece goods with unobserved cost functions,” 59 *Journal of Law and Economics*, 293-311.
- HU, XIAOXIAO AND HAORAN LEI. 2023. “Information transmission in monopolistic credence goods markets,” arxiv.org/abs/2303.13295.
- INDERST, ROMAN AND MARCO OTTAVIANI. 2012. “Financial Advice,” 50 *Journal of Economic Literature*, 494-512.
- MILGROM, PAUL AND JOHN ROBERTS. 1992. “Economics, Organization and Management,” Englewood Cliffs, N.J.: Prentice-Hall.
- OI, WALTER. 1971. “A Disneyland Dilemma: Two-Part Tariffs for a Mickey Mouse Monopoly,” 85 *Quarterly Journal of Economics*, 77-96.
- REIFNER, UDO, DORIS NEUBERGER, ROGER RISSI, CHRISTINE RIEFA, MICHAEL KNOBLOCH, SEBASTIEN CLERC-RENAUD, AND CHRISTIAN FINGER. 2013. “Study on the remuneration structures of financial services intermediaries and conflicts of interest,” Hamburg: Institut für Finanzdienstleistungen.
- TEH, TAT-HOW AND JULIAN WRIGHT. 2022. “Intermediation and Steering: Competition in Prices and Commissions,” 14 *American Economic Journal: Microeconomics*, 281-321.




Figure 1: Products traded in equilibrium. In the vertically shaded area the expert fully reveals meaning that all three products can be traded; in the diagonally shaded area she partially reveals, i.e., products α and β can be traded; in the horizontally shaded area she doesn't reveal so that only product α is traded.