

WORKING PAPER · NO. 2022-108

Superiority-Seeking and the Preference for Exclusion

Alex Imas and Kristóf Madarász

AUGUST 2022

Superiority-Seeking and the Preference for Exclusion

Alex Imas (Chicago Booth) and Kristóf Madarász (LSE)*

August 2, 2022

Abstract

We propose that a person’s desire to consume an object or possess an attribute increases in how much others want but cannot have it. We term this motive *superiority-seeking*, and show that it generates preferences for exclusion that help explain a host of market anomalies and make novel predictions in a variety of domains. In bilateral exchange, there is a reluctance to trade and people exhibit a ‘social endowment effect.’ People’s value of consuming a good increases in its scarcity, which generates a motive for firms and organizations to cater to such preferences by engaging in exclusionary policies. Randomly barring potential consumers from the opportunity to acquire a product increases the seller’s profits both in standard monopoly and auction settings. Such non-price-based exclusion leads to higher revenues than the classic optimal sales mechanisms. A series of experiments provides direct support for these predictions. In basic exchange, a person’s willingness to pay for a good increases as more people are explicitly barred from the opportunity to acquire it. In auctions, randomly excluding people from the opportunity to bid substantially increases bids amongst those who retain this option. Exclusion leads to bigger gains in expected revenue than increasing competition through inclusion. Our model of superiority-seeking generates ‘Veblen effects,’ rationalizes attitudes against redistribution, and provides a novel motive for social stratification and discrimination.

Keywords: social preferences, ownership, pricing, exclusivity, marketing, political economy, inequality, stratification, discrimination

*We are grateful to Nageeb Ali, Sullivan Anderson, Roland Benabou, Sylvain Chassang, James Choi, Stefano DellaVigna, Dan Gilbert, Shengwu Li, George Loewenstein, Ryan Oprea, Pietro Ortolova, Wolfgang Pesendorfer, Klara Steinitz, Adam Szeidl, Leeat Yariv for discussions and the Editor and four anonymous referees for excellent comments. We also thank seminar audiences at Arizona, Ashoka, Berkeley, Chicago, CEU, CMU, Harvard, Hebrew U, Hong Kong, LSE, Princeton, Vienna, Zurich, Peking University Business School, VIBES June 8, 2020, and the 2019 Warwick-Princeton-Utah Conference on Political Economy in Venice. All errors are our own. An earlier version of the paper was titled *Mimetic Dominance and the Economics of Exclusion*. Contact: alex.imas@chicagobooth.edu or k.p.madarasz@lse.ac.uk

1 Introduction

“Of all the passions, the passion for the Inner Ring is most skillful in making a man who is not yet a very bad man do very bad things. But your genuine Inner Ring exists for exclusion. There’d be no fun if there were no outsiders. The invisible line would have no meaning unless most people were on the wrong side of it. Exclusion is no accident; it is the essence.” — C.S. Lewis, *The Inner Ring*, 1944.

“The scarcity of the music not only makes the music itself enjoyable but also gives the collector a strange sense of superiority.” — Henry Rollins.

The desire to consume objects or possess attributes that others want but cannot have seems to be a significant driver of demand in a variety of settings. Voters appear to favor politicians who enact exclusionary policies which bar minorities and non-citizens from access to institutions and markets—in spite of the economic harm to themselves (see, e.g., [Mutz \(2018\)](#)). Firms and other entities appear to respond to a similar motive by practicing exclusionary policies that limit opportunities to purchase their goods or accessing their services. Sellers advertise widely, but maintain product shortages despite persistent excess demand at the going price.¹ Well-known restaurants and entertainment venues do not increase the price or the capacity despite long lines, and luxury brands would rather burn millions in undamaged product than threaten their exclusivity.² Moreover, such ‘artificial’ exclusivity and limited access often appears to be a feature rather than a bug, intended to drive up demand by highlighting the privilege of consumption in light of how much others would enjoy the same product, but cannot ([Cialdini, 2007](#)). While separate explanations have been proposed for some of these seemingly-disparate phenomena, we argue that they jointly point towards a model of preferences that generates a demand for exclusion.

In this paper, we argue that one’s desire for an object is not autonomous but intimately linked to the desires of others through a form of superiority-seeking, where people derive pleasure from consuming a good or possessing an attribute that others want more but cannot have. The prospect of exclusion then spurs demand by increasing the chance that those barred from acquiring an object want it more than those who end

¹For example, well-known street wear brand Supreme—which was valued at \$2.1 billion in 2020—sells branded, but otherwise fairly standard, clothing in limited quantities, charging massive markups over slightly lower quality products ([Von Wilpert, 2019](#)). Minimally horizontally differentiated, limited-edition variants of well-known products such as sneakers are marked up 800-1000% over similar versions ([Moreno, 2020](#)). Crucially, despite the price premiums, such exclusive products face considerable excess demand at the going price: new offerings typically sell out within minutes—with many missing the opportunity by mere seconds.

²See fashion brands destroying tons of unused product in order to maintain exclusivity, for example [Lieber \(2018\)](#).

up getting it. To build intuition, consider a fable on the rivalry between two siblings Luluwa and Awan. No matter how many toys they have at home, there is always one that one of the children wants to play with—the toy that the other child is playing with at the time. She wants to take it away and play with it only herself; the ability to exclude the other from playing with that toy makes it all the more desirable.³

We present a model of such preferences and empirically demonstrate their influence in several economically-important environments, such as basic exchange, classic monopoly, and auctions. We also outline implications for exclusionary practices more broadly, such as in luxury consumption, trade, redistribution, and discrimination. Four experiments provide direct support for the model while ruling out a host of alternative explanations.

In our model, person i 's utility from consumption is the sum of her private consumption utility (intrinsic taste) and a comparative term representing a imitative superiority-seeking motive. In our baseline characterization—which is generalized in the Online Appendix 1.1—this term corresponds to the largest gain in consumption utility that another individual in i 's social context would experience if he had her consumption in addition to his. This term is imitative (or mimetic) because it is defined over the desires of others; it captures superiority-seeking because this boost in utility requires these desires to be in excess of one's own and unmet.

We then derive the economic implications of superiority-seeking preferences. In bilateral trade, because owners value the unmet desire of potential buyers, there is a reluctance to trade due to a seller's lesser willingness to accept than a buyer's willingness to pay—a 'social' endowment effect (Kahneman et al., 1990). Our model also implies that people have a direct preference for objects that become relatively more scarce which generates a demand for both price- and non-price-based exclusion that can be exploited by firms as a tool for rent-seeking.

This motive has direct implications for markets. Our main result here considers the central monopoly problem: a seller producing identical copies of a good at a constant marginal cost to a pool of ex-ante identical consumers. The classic result shows that the monopolist maximizes profits by posting a price and allowing all who are willing to pay the price to buy. Instead, when consumers exhibit superiority-seeking, the seller always gains by randomly excluding a limited fraction of potential consumers and offering the product to the rest at a uniform price; artificially restricting supply creates excess demand at the going price, which boosts demand. We further show that engaging in such non-priced-based exclusion still dominates the classic optimal mechanism—even if random exclusion binds with heterogeneous probabilities across buyer types or is effective only for consumers with types lower than some threshold.

³According to some traditions, Luluwa and Awan were the partners and sisters of Cain and Abel.

We then consider competitive exchange. We show that randomly excluding potential buyers from the opportunity to bid in a first-price auction will lead to both higher bids and higher revenue for the seller of a private good despite the lower direct competition. Holding the level of exclusion constant, a larger number of active bidders still leads to higher average bids and revenue—the standard competitive force is intact. However, holding the number of active bidders constant, excluding a larger number of bidders also leads to higher bids and revenue. This is due to the increased chance that an excluded person may have greater desire for the item than the prospective winner, which boosts the latter’s valuation. This psychological effect is often stronger than the classic competitive effect such that excluding a positive fraction of bidders increases the seller’s expected revenue. If the degree of superiority-seeking is sufficiently large, randomly excluding bidders again generates higher rents for the seller than employing any reserve price under full inclusion—the classic optimal selling mechanism (Myerson, 1981; Riley and Samuelson, 1981).

Four experimental studies test the predictions of the model. The first considers the simple non-competitive setting of basic exchange. Participants were incentivized to report their maximum willingness to pay (WTP) for a unique good—a custom T-shirt designed specifically for the experiment—in one of two treatments. In the Baseline treatment, all participants could submit their bid for the good. In the Random Exclusion treatment, the experimenter publicly announced that a subset of people would be randomly barred from the opportunity to purchase the good, allowing only the remaining to submit their bids. Despite its transparent nature, exclusion led to a nearly 50% increase in willingness to pay amongst those who retained this opportunity.

The second study considers the competitive setting of first-price auctions where a single good is allocated to the highest bidder. Participants reported their bids for the same unique good as the first study in one of three treatments. In the Baseline treatment, all potential bidders participated in the auction. In the Random Exclusion treatment, a subset of subjects were randomly and publicly barred from the opportunity to participate. In the Non-Random Exclusion treatment, before announcing the auction, participants first reported the extent to which they wanted the object. In this treatment, exclusion *per se* was the same as in the Random Exclusion treatment, but those who retained the opportunity to bid knew that the desire of the excluded was lower than their own.

Three main results obtain. First, findings from our Baseline treatment confirm the standard prediction that bidding becomes more aggressive as the group size increases. Second, consistent with superiority-seeking, random exclusion spurred more aggressive bidding *and* increased the seller’s expected revenue relative to the Baseline condition. Strikingly, the increase in average bids due to exclusion is nearly double the impact

of increasing competition. For example, the average bids from a group of *four* active bidders with exclusion—where two are randomly excluded from a group of six ($M = 6$ and $K = 2$)—were nearly 60% higher than the average bids from a group of *six* active bidders without exclusion ($M = 6$ and $K = 0$). Holding the number of active bidders constant, we also find that the presence of random exclusion substantially increases bids and expected revenue. Finally, consistent with the model’s prediction, we find that both average bids and expected revenue in the Non-Random Exclusion treatment were similar to those in the Baseline treatment and lower than under Random Exclusion.

Studies 3 and 4 demonstrate the robustness of the results and even more directly test the implications of the model. Both pre-registered studies were conducted online in a fully anonymous setting. An online store was set up to deliver a unique good—a print of a painting by one of the authors—to participants. Study 3 matched an Active participant with three Passive participants to form a group of four. After reporting her desire for the good, the Active participant was incentivized to report her maximum WTP for it in each of three scenarios. The scenarios differed only in the number of Passive participants who would potentially be excluded from the opportunity to purchase the good. Moreover, Active participants were randomly assigned to a High, Low, or No Information treatments. In the High (Low) Information treatment, the Active participant was told that the Passive participants in her group had a greater (lower) desire for the good than her own; in the No Information treatment, Active participants were not given any information about the desires of the group. This experiment allowed us to explore superiority-seeking on the intensive margin, in a within-subject design where the independence between exclusion and actual scarcity is readily transparent. Consistent with our framework, Active participants reported higher WTP as the level of exclusion increased—but only in the High and No information treatments.

Study 4 was designed to replicate the classic monopoly setting. We recruited 100 participants and elicited their valuation for the same art print as in Study 3 across two scenarios. In the No Exclusion scenario, all participants would have the opportunity to purchase the art print, depending on their willingness to pay; in the Exclusion scenario, only sixty of the recruited participants would be randomly selected to have this opportunity. All participants were told that after they had reported their willingness to pay in both scenarios, the computer would flip a coin to determine which would be played out ‘for real.’ Results showed that exclusion—which was transparently random by design—shifted demand substantially to the right: median WTP increased by 50% in the scenario where participants knew that others would be excluded. Moreover, as long as the marginal cost of supplying the product is not too low (at least \$2 in our setting), a profit-maximizing firm would be better off in the exclusion scenario despite the 40% smaller customer base. Together, our four studies rule out a host of alternative

explanations such as scarcity as a signal of value, direct consumption externalities, interdependent values *per se*, e.g., [Milgrom and Weber \(1982\)](#), social preferences over money, or a heightened desire to obtain the good due to differential attentional effects, or feeling lucky due to not being excluded.

While to the best of our knowledge the superiority-seeking motive has not been considered in economics, the idea that it is an important aspect of social interactions has a long tradition in social thought, e.g., [Augustine \(2009\)](#), [Hobbes \(1998\)](#), and [Rousseau \(1755\)](#). For example, in his *Discourse on the Origin of Inequality*, [Rousseau](#) emphasizes the role of what he considers to be the critical and potentially destructive motive of human sociability: a person’s tendency to compare herself to others and engage in activities whereby she can experience her superiority over them. Such *amour propre*, as Rousseau calls it, represents a person’s concern with comparative success or failure as a social being and involves joy from feeling superior over others; he describes the motive as key for understanding the function of political institutions.⁴ More recently, Dr. Martin Luther King referred to a similar motive as a universal ‘Drum Major Instinct.’⁵

We argue that economic settings naturally spur superiority-seeking over the unmet desires of others. In Section 6, we outline broader implications of our framework for both non-price-based and priced-based methods of exclusion. In the presence of binding income inequality, superiority-seeking predicts so-called ‘Veblen effects’ ([Veblen, 1899](#)), where the demand for luxury goods increases in response to price increases. This seeming violation of the law of demand has typically been explained through the motive of signaling one’s income status.⁶ Our model provides an alternative and complementary private mechanism which operates irrespective of the direct observability of one’s own consumption or income. To illustrate, consider the example of a real diamond ring versus a fake one (e.g., cubic zirconia); both are practically indistinguishable to outside observers. Under signalling, it may thus make sense to buy the cheap fake ring and use the savings for some other form of consumption. In our framework, the accompanying superiority boost is only present for the true diamond, as others’ unmet desire for a fake diamond is likely close to null and burning money delivers no benefits. As we discuss, the predictions of superiority-seeking for generating Veblen effects closely match key aspects of the consumer landscape for luxury goods and help rationalize various aspects that a pure signalling account may struggle to account for.

⁴Our formulation also draws on the work of the literary scholar René [Girard \(1965, 2004\)](#) who emphasizes a distinction between appetites—consumption utility in our context—and more intense desires. In his language, appetites are basic and individual. Desires, on the other hand, are not autonomous but inherently mimetic.

⁵Dr. King’s sermon can be found here: <https://kinginstitute.stanford.edu/king-papers/documents/drum-major-instinct-sermon-delivered-ebenezer-baptist-church>.

⁶See, for example, [Bagwell and Bernheim \(1996\)](#); [Pesendorfer \(1995\)](#).

The use of non-price-based methods of artificially restricting availability of products and services is common in a variety of domains. Advertising often exploits the psychology of exclusion in sales of private goods such as in the case of scarcity marketing—a cornerstone of marketing practice.⁷ Many advertising guides explicitly note that the practice of scarcity marketing rests on the premise that access to exclusive goods makes owners ‘feel powerful’ as a result of obtaining something that others desire but cannot have.⁸ The prevalence of ‘club goods’—which are characterized by exclusion as a feature rather than a bug, despite these being non-rivalrous on the margin—has also long puzzled economists. Becker (1991) notes the lack of firm responses to excess demand, highlighting the puzzling presence of long queues for products such as restaurants, nightclubs, and sporting events. He considers several explanations but ultimately concludes that they would not generate the unwillingness of firms to either expand supply or raise the price. Our model can rationalize these phenomena. Firms and advertisers may artificially restrict actual or perceived availability of a good so that those who end up attaining it get a utility boost from the unmet desires of others. As our fourth study demonstrates, the gap between supply and demand can be an effective tool to extract rents from consumers, such that eliminating excess demand may actually lead to a *drop* in overall demand.

More broadly, superiority-seeking provides a cohesive explanation for exclusion by institutions, as well as for group-based discrimination. ‘Included’ members derive greater pleasure from consuming goods, possessing attributes, and belonging to organizations from which others are restricted—even if these restrictions are unnecessary from a marginal cost perspective (e.g., in case of ‘club goods’), lead to material losses (as in the case of immigration restrictions and protectionism), or are based on seemingly arbitrary criteria or characteristics (such as in the case of ‘taste-based’ discrimination). Finally, we discuss how superiority-seeking rationalizes seemingly anomalous attitudes towards redistributive policies, such as the observation that the strongest opposition to increases in the minimum wage is concentrated amongst those making the second-to-lowest amount (Kuziemko et al., 2014), and provides a psychological basis for the maintenance of social stratification (Darity Jr et al., 2015).

⁷In his runaway bestseller on the topic, titled *Influence: The Psychology of Persuasion*, Robert Cialdini identifies scarcity marketing as one of the six ‘Weapons of Influence.’ He argues that scarcity marketing is most effective when people are made to believe that rivalrous others are simultaneously competing for the same good—that attaining it means others who want the good are excluded from having it.

⁸See practical guides to marketing: <https://sumo.com/stories/scarcity-marketing>, <https://blog.crobox.com/article/scarcity>

2 A Model of Superiority-Seeking

2.1 Setup

In this Section, we develop our basic model of imitative superiority-seeking and outline some of its consequences. Our model is purposefully simple. As we describe below, we present a more general model in Online Appendix 1.1 where we relax various simplifying assumptions and demonstrate the robustness of our results.

To illustrate, suppose first that there is a single object and two people i and j . Let their consumption utilities (intrinsic tastes) from the object be v_i and v_j , respectively. Person i 's overall utility from consuming the good is then given by:

$$\text{person } i\text{'s valuation} = \underbrace{v_i}_{\text{consumption utility}} + \alpha \underbrace{\max_{m \in \{i,j\}} \{v_m - v_i\}}_{\text{superiority boost}}$$

where $\alpha \in [0, 1)$ is the strength of the superiority-seeking motive.⁹ If $\alpha = 0$, there is only classic consumption utility. If $\alpha > 0$, there is the additional presence of the superiority boost: person i enjoys a utility boost from consuming the good proportional to the extent to which j would derive greater consumption utility from the object than she does. If i has a greater intrinsic taste for the object, the comparative term is zero; otherwise, it is positive. Person i 's overall utility from consuming the object thus mirrors j 's unmet intrinsic taste for it as long as this is in excess of her own.

Consider now a more general setting. Let there be M people. Let each person i 's consumption be given by a non-negative L -dimensional vector, $c_i \in C_i \subseteq \mathbb{R}^L$, where each dimension corresponds to an attribute or type of good. Direct consumption utility is given by $V_i(c_i, t_i) = \sum_l v_{i,l}(c_{i,l}) + t_i$, where $t_i \in \mathbb{R}$ is monetary transfer. We thus assume that preferences are additively separable across dimensions and are quasi-linear in money.¹⁰ Finally, each $v_{i,l}$ is bounded and increasing with $v_{i,l}(0) = 0$. Let C then be the product of the individual C_i sets with generic element $c \in C$. Person i 's utility, $U_i(c, t_i) : C \times \mathbb{R} \rightarrow \mathbb{R}$, is given by:

$$V_i(c_i, t_i) + \alpha \sum_l \max_{m \in M \setminus i} \max \{v_{m,l}(c_{m,l} + c_{i,l}) - v_{m,l}(c_{m,l}) - v_{i,l}(c_{i,l}), 0\}.$$

In words, in addition to her standard consumption utility, a person i derives utility from superiority-seeking. For each kind of good or attribute, this superiority boost

⁹The parameter α is constrained to be less than one in our setting to ensure that a person's overall utility is *non-decreasing* in her own consumption.

¹⁰Considering an additively-separable specification of consumption utility is common in leading behavioral models of consumption with a comparison point, e.g., [Kőszegi and Rabin \(2006\)](#), and [Bordalo et al. \(2013\)](#), and in tractable choice, e.g., [Friedman and Sákovics \(2015\)](#) or [Camara \(2021\)](#). One can, however, extend the model to consumption of a given kind being a vector rather than a scalar.

corresponds to the maximal excess consumption utility *gain* that another person j would derive if he had, in addition to his own consumption of the good, i 's consumption of the good transferred to him as well. If i has the greatest direct utility from this consumption, then this boost is zero; otherwise it is positive. Our setup implies no superiority-seeking over money per se; rather, it is directed towards objects of consumption. A few remarks are in order.

Remark 1. In our formulation, for each attribute, the extent of superiority associated with one's consumption is determined by the *maximal* other's excess taste for one's consumption of this attribute. This assumption is somewhat extreme and is adopted for simplicity only. In particular, it implies that holding the maximum excess taste for i 's consumption constant, it does not matter what fraction of others have an excess taste for what i has, nor does it matter how much those others would like to have what i has.

In Online Appendix 1.1 we generalize the model. We relax the maximum specification and allow the superiority boost to be determined by the convex combination of the *average* excess tastes and the maximum excess taste of others. For any such combination, the utility boost now depends positively on the *fraction* of others who would derive an excess consumption utility from i 's consumption and is strictly increasing in each other person's excess consumption utility from it—even when the maximal excess valuation is unchanged. Furthermore, since the combination is always smaller than the max, the impact of superiority-seeking relative to consumption utility, given any endowment c , is now always quantitatively *smaller* than in the above formulation. Nevertheless, we show that all our predictions extend either to all such generalizations considered, or for weighted averages well away from the max specification. ¹¹

Remark 2. The formulation above also assumes that the utility boost enters as an additive term next to consumption utility. In Online Appendix 1.1, we also generalize this aspect of the model and allow it to be a multiplicative factor of consumption utility; there, if one derives no utility from the consumption of an attribute, she also

¹¹Formally, for a given c , $j \in M \setminus i$, and fixed l , where for ease of exposition we now suppress l from the notation, let $v_{j,i} \equiv \max\{v_j(c_j + c_i) - v_j(c_j) - v_i(c_i), 0\}$ denote j 's excess valuation of i 's consumption. Person i 's utility is now given by:

$$v_i(c_i) + \alpha(1 - \beta) \frac{\sum_{j \in M \setminus i} v_{j,i}}{M - 1} + \alpha\beta \max_{j \in M \setminus i} v_{j,i} + t_i, \quad (1)$$

where $\beta \in [0, 1]$. If $\beta = 0$, the boost corresponds to the average excess valuation of others. If $\beta \rightarrow 1$, it converges to the maximum specification described in the main text above. For *any* given $\beta \in [0, 1]$ and c , the weight assigned to the excess valuation of each $j \in M \setminus i$ is strictly positive, and the weights always add up to one. As β increases, the boost also increases. At the same time, i 's overall utility continues to be increasing in her own consumption c_i and in the excess valuation of any j for any $\alpha \in [0, 1)$. If a prediction described below holds only for β sufficiently high, we point this out. Otherwise it holds for all $\beta \in [0, 1]$.

derives no superiority boost from possessing this attribute even if others have positive unmet intrinsic taste for it. Hence, there is now a non-monotonic relationship between consumption utility and the superiority boost. We again show that all our main predictions are robust to such multiplicative specifications as well.

Remark 3. Superiority-seeking does not simply correspond to imitative or ‘mimetic’ desire, or interdependent values, e.g., [Milgrom and Weber \(1982\)](#), or preference uncertainty/common values (e.g., [Fischer et al. \(2000\)](#), [Li and Mattsson \(1995\)](#), [Wilson \(1969\)](#)), whereby one infers the quality or the consumption value of a good from how much others like it, irrespective of whether or not others have the good. Crucially, the extra utility from consuming an object in our setup does not derive from the fact that many like such a good as well; it is not a matter of interdependent values *per se*. Rather, the utility boost is derived from the unmet desires of others—from what they ‘lack.’ It corresponds to a motive whereby people enjoy consuming goods through exclusion and excess.

Superiority-seeking is also distinct from the idea of envy as is commonly expressed in the context of envy-free allocations, e.g., [Varian \(1976\)](#). There, envy refers to person i preferring what j has over her own allocation. It relates to what person i desires as a function of what person j has. The motive of superiority-seeking instead refers to what the *other*, person j , desires in relation to what person i has.

Remark 4. A primitive of our model is the comparison set M . In some applications, the comparison set arises naturally given the set of people one interacts with: one’s set of siblings, high school class, set of business partners or colleagues, sets of peers, or known social media contacts, etc. In line with leading approaches to social preferences, we do not endogenize the comparison set, but implicitly assume that such preferences are framed somewhat more narrowly and need not be global.¹² We also do not claim that superiority-seeking applies to all goods and attributes equally—it may be affected by salience ([Bordalo et al., 2013](#)), and there may be factors that mitigate its impact or shift it from one domain to another. Our formulation potentially allows for this type of heterogeneity and dependence. Understanding these factors is clearly important. We are unaware of any prior empirical research that would help guide these modelling choices and leave such extensions to future work.

¹²The social context that defines a comparison or consideration set is a basic assumption in nearly all models of social preferences, e.g., [Fehr and Schmidt \(1999\)](#) and [Charness and Rabin \(2002\)](#), and the empirical results that motivated these models therein can only be interpreted given such context-dependent bracketing. While this is a limitation of our framework as well, Sections 6.2 and 7 discuss how comparison sets may be shaped in the context of political economy and advertising.

2.2 Trade

We now turn to basic implications of imitative superiority-seeking for bilateral trade. Let there be two people and a single good. Suppose that each party's consumption utility is drawn i.i.d. from a strictly increasing cdf $F(v)$ which admits a continuous density defined on some bounded interval $[0, \bar{v}]$. The good is allocated randomly to one of the parties.

First, let there be no private information about preferences; i.e., the realizations of the consumption utilities become public before trade. Suppose that any monetary transfer is possible and the parties bargain efficiently. Our first corollary shows that superiority-seeking is a force against trade.

Corollary 1. *If $\alpha = 0$, trade happens with an ex-ante probability of $1/2$. If $\alpha = 1$, trade never happens. For any transaction cost $\varepsilon > 0$, the ex-ante probability of trade is strictly decreasing in α .*

The logic of why trade never happens with full superiority-seeking is simple. If the owner were to sell the object to the buyer, then whatever is the buyer's gain from the trade in terms of consumption utility, this is mirrored directly as the seller's loss of imitative superiority. Trade is effectively zero sum. More generally, if person j were to gain more in consumption utility than i would lose—the very precondition of trade with money— α fraction of this would correspond to a psychological loss for i . Given any positive transaction cost, the ex-ante probability of trade then strictly decreases in superiority-seeking and becomes zero before its extent becomes full.

Suppose now that valuations are realized privately and consider a simple exchange mechanism. Consider any price p from $(0, \bar{v})$ and suppose that players simultaneously decide whether to say yes or no to trade at this price. These decisions are then publicly announced and trade takes place at p if and only if both parties said yes. We focus on BNE with a positive ex-ante probability of each player saying yes.¹³

Corollary 2 (Social Endowment Effect). *If $\alpha = 0$, the probability that the seller says yes is the same as the probability that the buyer says no. If $\alpha > 0$, the former is strictly lower than the latter for any given price p .*

The logic of the above result is strategic. Given superiority-seeking, if the buyer values getting the item more, the seller values keeping it more. The seller is thus reluctant to say yes because conditional on trade, her valuation from keeping the object

¹³This restriction is imposed only because there is always another equilibrium where both the seller and the buyer announce no. That equilibrium is, however, purely artificial and is not robust to any trembles.

may increase. Hence, in equilibrium the seller needs to be compensated for the loss of imitative superiority.¹⁴

The above is consistent with the classic finding of an endowment effect (Kahneman et al., 1990). What drives the endowment effect here, however, is social and relates to a person’s belief about the preferences of others. In turn, the trading mechanism is also key. For example, consider the case of one-sided forced trade whereby for a given price p , one of the parties can unilaterally decide whether or not to trade. There, the probability that for any given price the seller says yes matches the probability with which the buyer says no to trade. This is in contrast with a loss-aversion explanation that would still predict a wedge between the seller’s WTA and the buyer’s WTP (Thaler 1980; Köszegi and Rabin 2006, with the reference point at no expectation of trade). Note that we do not claim that some form of reference dependence does not play a role in generating observed WTA-WTP gaps; rather, we propose superiority-seeking as a distinct, complementary channel with distinct comparative static predictions.¹⁵

More generally, in our setup a person’s willingness to trade an object is a function not only of the identity and the endowment of the recipient, but the endowments of those unaffected by trade. This leads directly to the implications of superiority-seeking for people’s preferences over scarce goods and exclusion.

2.3 Scarcity and Exclusion

Superiority-seeking leads to an increased willingness to keep goods that become relatively more scarce. To illustrate this, consider a setting where P randomly chosen people are assigned a pen and C other randomly chosen people are assigned a cup, with $P + C < M$. Suppose that each person has some privately known unit demand for a pen and separately for a mug drawn from a non-degenerate distribution. A single randomly chosen person i has the right to unilaterally swap her object with a randomly chosen other who is assigned the opposite object.

Corollary 3. *If i is a cup owner, the probability that she swaps is strictly decreasing in P and strictly increasing in C .*

¹⁴Note that this mechanism may also help rationalize the so-called dynamic ‘reactive devaluation’ effect in negotiations (Ross, 1995).

¹⁵For example, we would predict a smaller WTA-WTP gap in cases where the seller cannot influence the buyer’s access to the good through her WTA. To that end, Plott and Zeiler (2005) amended the classic endowment effect paradigm by making it salient that buyers and sellers were trading with the experimenter—who presumably has a large quantity of the good—rather than with one another. This was done by placing the good on the desk of both the buyer and the seller, rather than just the seller. They found almost no WTA-WTP gap. Similarly, we would predict that the endowment effect is attenuated when buyer and seller valuations are elicited separately, such that both are trading with the experimenter. This is consistent with the findings of Shogren et al. (1994). We would like to thank an anonymous referee for this suggestion.

The logic of the above prediction is based on the fact that as the relative scarcity of a good increases (decreases), the superiority boost associated with keeping it increases (decreases). All else equal, the more scarce an object is, the greater is the expected superiority boost associated with consuming it since the expected excess valuation for this object increases as well.

An analogous way to express this preference for exclusivity is to consider consumption along a single dimension. As long as consumption utility exhibits diminishing differences, it follows that, all else equal, the value of one's consumption is higher the lower the consumption of others along this dimension is. This also implies that a person may also care about *who* she trades with under a given fixed price. To illustrate, suppose that person i has two Swiss watches, and everyone else in her social context is randomly endowed with either one or zero such watches. All else equal, person i 's willingness to give up one of her Swiss watches in exchange for the fixed price will be lower if it is to someone who already has a Swiss watch than if it is to someone without one. In expectation, the former trade preserves a greater superiority boost than the latter.

2.4 Monopoly

Let us now turn to our central result of this Section by considering the classic monopoly problem. The seller can produce identical copies of a good at some constant marginal cost normalized to zero. Each buyer i has a unit demand for the good. We show that the seller can always achieve a higher revenue by randomly excluding some buyers from the opportunity to acquire the good—provided there are enough buyers left—than by allowing all to buy at the classic optimal monopoly price. Creating excess demand becomes an effective tool for rent seeking.

Before demonstrating this, let us first describe some implications in a simple example under the (unrealistic) assumption of consumption utilities being *public* information. Under standard preferences, the seller wants to sell to each buyer at a price equal to her reservation price (perfect price discrimination). In the presence of superiority-seeking, the seller can instead maximize her revenue by excluding buyers who want the item most.

Example 1. *Suppose there are three people with $v_1 = v_2 = l < h = v_3$.*

1. *If $\alpha < h/2(h-l)$, the seller's revenue is maximal when selling to each buyer i at $p_i = v_i$.*
2. *If $\alpha > h/2(h-l)$, the seller's optimal revenue is given by excluding the high value buyer and selling to the low value buyers at $p = l + \alpha(h-l)$.*

While the seller never gains by excluding the lowest valuation type, if the extent of superiority-seeking is sufficiently high, she gains from excluding the high valuation buyer. Such exclusion generates a boost in aggregate demand larger than the high-valuation buyer's maximal willingness to pay.¹⁶ More generally, in this setting exclusion can benefit the seller even if the seller cannot exclude the highest type but only some medium types.

Consider now the standard monopoly setting. Each buyer's consumption utility is again drawn independently from some strictly increasing cdf F which admits a continuous pdf and is her private information. The classic result, given standard preferences ($\alpha = 0$), is that the seller's optimal selling mechanism is to set a single price common to all potential customers and to allow each to buy at will (e.g., [Harris and Raviv 1981](#)). Our next result implies that for any $\alpha > 0$, if the number of potential buyers is not too small, randomly excluding some buyers from the opportunity to buy the product always permits the seller to achieve a strictly higher revenue.

Proposition 1. *For $\alpha = 0$, the seller never gains from randomly excluding a buyer and her optimal profit is achieved by posting some price p^* to all. For any $\alpha > 0$, there exists M_α such that if $M > M_\alpha$, the seller's revenue is strictly higher when she randomly excludes a strictly positive number of buyers and sets a common price p' for the rest.*¹⁷

The above result helps parsimoniously resolve a key puzzle in pricing practices, e.g., [Becker \(1991\)](#), with minimal assumptions on the underlying value distribution and without imposing artificial production constraints. In contrast to the classic optimal mechanism where the uninformed seller sets the optimal monopoly price to all and allows each to buy at will, in the presence of superiority-seeking, the seller always gains by engaging in non-price-based *exclusion* and creating excess demand for her product at the going uniform price. To see the logic, note first that in the classic case, random exclusion strictly reduces demand without allowing the seller to extract a higher rent from the remaining buyers. Hence, she can only lose from such non-price based exclusion. In the presence of superiority-seeking, however, the non-excluded buyers now derive extra pleasure from superiority boost when acquiring the product. This outweighs the loss from having a smaller consumer base if there are enough consumers

¹⁶The above also illustrates that superiority-seeking leads to very different conclusions than a resale motive. While the first point holds regardless of whether or not resale between the buyers is possible, the second point holds only if resale is sufficiently costly. If resale was free, the low valuation buyers would want to sell their objects to the high valuation buyer and anticipating this, their total willingness to pay in equilibrium would need to be bounded from above by $l + h < 2l + h$.

¹⁷In the generalization of our model, if random exclusion strictly benefits the seller for a given M and α , this holds also for any $\beta < 1$ that is not too low. Furthermore, holding such an M fixed, an increase in α increases the range of β 's where this holds as the superiority-seeking term smoothly increases both in α and in β .

left.

Above, exclusion was random per buyer and we assumed that random exclusion was binding irrespective of the excluded buyers' intrinsic tastes. In practice this assumption may be too strong; in some contexts, buyers who really like the product may still find a way of obtaining the rationed good and this may be anticipated by others.¹⁸ However, assuming that random exclusion binds for all types equally is not needed for this result. First, it follows from the proof of Proposition 1 that the same result holds if exclusion binds only probabilistically, whereby there is some strictly positive, but potentially heterogeneous, probability that exclusion binds for any given type of a randomly targeted buyer. If it does not bind, this buyer can still obtain the product at the going price. Second, one can also consider v^h -constrained random exclusion whereby random exclusion binds if and only if the excluded buyer's consumption utility is *lower* than the cutoff v^h ; otherwise, the buyer is still able to obtain the good at the charged price. The next corollary demonstrates the robustness of Proposition 1 to this extension as well.

Corollary 4. *For any $v^h > 0$, Proposition 1 extends when replacing random exclusion with v^h -constrained random exclusion, either weakly or strictly (always strictly if $v^h > p^*$).*

Finally, note that the increase in the willingness to pay by the non-excluded consumers is increasing in the number of those excluded. While the optimal fraction of potential buyers to be randomly excluded depends on further details about the cdf F , the next example illustrates that this fraction may be substantial.

Example 2. *Suppose F assigns probability $1/3$ to $v_i = 1$ and $2/3$ to $v_i = 0$. Let $M = 18$. The optimal number of buyers to randomly exclude is 5 for any $\alpha > 0.53$ and zero otherwise.*

3 Exclusion in Auctions

Consider now the implications for the classic competitive allocation mechanism of a first-price auction. Potential buyers compete for an indivisible good by submitting sealed-bids. As before, each person i 's private consumption utility for this object v_i is an independent and private draw from a cdf $F(v)$ over $[0, \bar{v}]$ with a bounded, strictly positive, and continuous density. The seller derives a normalized consumption utility of zero from the object. Each bidder then maximizes her expected utility, which is given by

$$E[v_i + \alpha \max_{j \in M} \{v_j - v_i\} - b_i] \text{ in case of winning, } 0 \text{ otherwise.}$$

¹⁸Note that the discussion here implies that for Proposition 1, it suffices that people *believe* that there are others with taste greater than their own who cannot obtain the good due to excess demand.

We denote the number of people randomly excluded from the ability to submit a bid at the start by $K < M - 1$, and the number of active bidders by $N = M - K \geq 2$. As is standard in the auction literature, we focus on the monotone symmetric equilibrium where the lowest type makes zero surplus (Milgrom and Weber, 1982).

Proposition 2. *The symmetric equilibrium is characterized by a bidding strategy $b(v_i)$ such that*

1. *If $K = 0$, bidding is independent of α .*
2. *If $K > 0$, bidding and revenue are increasing in α .*
3. *Holding K constant, bidding and revenue are increasing in M for any $\alpha \geq 0$.*
4. *Holding N constant, bidding and revenue are increasing in M iff $\alpha > 0$.¹⁹*

First, irrespective of the degree of superiority-seeking, the seller's expected revenue is increasing in the number of competing bidders. The classic competitive force is intact. Furthermore, in the absence of exclusion, bidding is independent of the superiority-seeking motive. At the same time, exclusion also leads to more aggressive bidding given any positive α and this effect increases in the number of those excluded.

To see the logic, note that in the absence of exclusion, the auction is efficient; the winner is the player with the highest realized valuation. Hence, conditional on winning, she derives no boost from superiority. In contrast, in the case of exclusion, the winner of the auction may not be the person with the highest valuation for the object. She then derives a utility boost from superiority since those randomly excluded may have excess valuation for the object. This leads to more aggressive bidding by those included.

From this, it follows that now both more inclusion (greater competition, higher N) and more exclusion (higher K)—which leads to lower competition—result in higher bids. It is therefore natural to ask whether it is better for the seller to have *more* or *less* buyer competition from a given pool. The answer depends on the extent of superiority-seeking. For simplicity, below we assume that private valuations are drawn uniformly and denote by $\Pi(M, K)$ the seller's expected revenue when K of the M potential bidders are randomly excluded at the start.

Proposition 3. *For any M , there exists $\alpha^* < \alpha_M < 1$ such that*

1. *If $\alpha < \alpha^*$, $\Pi(M, K)$ is decreasing in K ;*
2. *If $\alpha > \alpha^*$, $\Pi(M, K) > \Pi(M, 0)$ if $K \leq K_{M, \alpha}$ with $K_{M, \alpha}$ increasing in α and M ;*

¹⁹Points 1, 2 and 4 of Proposition 2 continue to hold for any β . In the case of point 3, for simplicity, we restrict attention to $K = 0$, where again the prediction extends to any β .

3. If $\alpha > \alpha_M$, $\Pi(M, K)$ is increasing in K ;
4. If $\alpha^* < \alpha < \alpha_M$, $\Pi(M, K)$ is inverse U -shaped in K .

If the extent of superiority-seeking is small, the standard result holds and more competition is better—exclusion hurts the seller. For moderate levels of superiority-seeking, the comparative static is U -shaped. Exclusion, despite the lower competition, raises the seller’s expected revenue up to a threshold.²⁰ If the number of excluded crosses this threshold, this effect reverses. This threshold is increasing in the degree of superiority-seeking and in the overall size of the group. Finally, if superiority-seeking is sufficiently high, the seller’s expected revenue is strictly increasing in the number of bidders excluded.

To see the logic, consider first the standard effect. Exclusion decreases bid amounts because each bidder knows that she faces less competition, and thus she has an incentive to shade more. Furthermore, there are now fewer bidders which further decreases the seller’s rents. In the presence of superiority-seeking, this standard effect is fully present. However, there is now a countervailing psychological force. Since the winner may not be the person with the highest consumption utility, the winner experiences an added superiority boost. As the number of excluded bidders increases, competition decreases, this countervailing psychological force also becomes stronger since the expected maximal unmet excess valuation also increases. If the extent of the superiority-seeking motive is not too small, initially, the force due to superiority-seeking outweighs the standard force due to lesser competition. As long as the superiority-seeking motive is not too high, this balance reverses as the number of excluded individuals reaches a threshold. Past this threshold, further exclusion decreases the seller’s revenue.

Crucially, random exclusion may not only lead to higher revenue than full inclusion, but, just as in the monopoly setting of Section 2, to a revenue that is even higher than from employing the classic optimal mechanism. Given independent private values and the standard regularity condition, $\alpha = 0$, the seller’s optimal revenue is achieved by the mechanism of running the first-price auction with an optimally-set reserve price (Myerson, 1981). The classic result on the value of competition further implies that under $\alpha = 0$, the revenue with $M + 1$ bidders, maintaining full inclusion and no reserve price, leads to a revenue greater than the seller’s optimal revenue with M bidders, achieved by employing the optimal reserve price and full inclusion, (Klemperer and Bulow, 1996). As the corollary below shows, if α is sufficiently high, then randomly ‘subtracting’ bidders from a pool of M potential bidders leads to even higher revenue than from the auction with $M + 1$ bidders maintaining full inclusion.

²⁰This reversal of the classic revenue prediction continues to hold under the generalization of our model as long as β and α are not too low.

Corollary 5. *Suppose that $M \geq 4$. If α is sufficiently high, then for any $K \geq 2$, it follows that $\Pi(M + 1, 0) < \Pi(M, K)$.*

The above discussion then implies that if the extent of superiority-seeking is sufficiently high, randomly excluding a sufficient number of bidders and then running the auction without a reserve price is not simply better than no exclusion with full competition, but it again dominates the classic optimal selling mechanism even in competitive exchange.

Lowest Exclusion. We now consider the case where the bidders who are excluded are not random, but commonly known to have intrinsic tastes that are *lower* than the tastes of those included. We return to the general setup from the beginning of this section, but suppose that it is the K lowest valuation bidders who are initially excluded from the opportunity to submit a bid (without disclosing any information about their realized valuations). We refer to this as ‘bidding under lowest exclusion’ and denote the seller’s expected revenue under such lowest exclusion by $\Pi_L(M, K)$.

Under lowest exclusion, standard revenue equivalence now follows irrespective of α : the seller’s expected revenue under lowest exclusion is the same as under no exclusion for any given M and feasible K . The bidding function is also independent of α : in both cases the winner of the auction is the person with the highest consumption utility, hence she derives no superiority boost conditional on winning, and the expectation of the second-highest consumption utility across the active bidders is the same. The only difference between lowest exclusion and no exclusion is that consumption utilities and thus realized bids are positively selected under lowest exclusion.

Proposition 4. *1. The bidding strategy under lowest exclusion is independent of α . Furthermore, $\Pi_L(M, K) = \Pi(M, 0)$ for any α , K , and M . 2. If $\alpha = 0$, the bidding strategy under lowest exclusion is higher than under random exclusion and $\Pi_L(M, K) > \Pi(M, K)$ for any given M and K .*

Importantly, if $\alpha = 0$, the average bid should be *higher* under lowest exclusion than under random exclusion. This is based on the fact that in the former case, the included bidders come from a positively selected pool—the mechanical effect—and that each active bidder also has a lower incentive to shade her bid—the strategic effect of positive selection. However, if α is not too low, superiority-seeking is predicted to reverse this relationship as per the previous discussion.

4 Experimental Studies

We now turn to directly testing the predictions of the model. Our first experiment is in the context of a classic exchange mechanism where everyone can potentially attain a

good and there is no direct competition between people. This study tests the predictions outlined in Section 2, and more specifically, Proposition 1. We then test the predictions in the context of the auction settings outlined in Section 3. Here, people compete for a single good. Together, the studies allow us to identify superiority-seeking and the proposed preference for exclusion while ruling out alternative explanations.

4.1 Study 1: Basic Exchange

Consider a setting with M people and M identical goods. Each buyer has a unit demand and consumption utilities are drawn exactly as in Section 2. Take a simple Becker-DeGroot-Marschak (BDM) selling mechanism. Let K denote the number of potential buyers randomly excluded from the opportunity to acquire the good. Subsequently, each non-excluded individual has to submit a non-negative bid. After the bids are submitted, a price is drawn randomly from $[0, \bar{v}]$ according to a positive and bounded density. An item is sold at this *common price* to each bidder whose bid is weakly higher than this amount and not to others.

Proposition 5. *Suppose that each person who submits a bid b_i buys at a randomly drawn price p if and only if $b_i \geq p$. In a symmetric equilibrium each player bids $b_i = v_i + \alpha E \max_{j \in K} \{v_j - v_i, 0\}$.*

In the absence of exclusion, the bids reveal the buyers’ private consumption utilities. In the presence of exclusion, the bids also reflect the influence of superiority-seeking. Critically, the above difference in bids between the case with and without exclusion depends on the excess desire of others being *unmet*. If superiority-seeking was purely about excess desire regardless of whether they are met or not, then bids would depend only on M and not on K . Similarly, if valuations for the object were simply interdependent in the classic sense, there would be no *ex-ante* expected difference in bids as a function of exclusion. This is true because there is no direct competition between bidders and the distribution of signals about valuations people receive, and hence a person’s expectation of the value of obtaining the object would have to be the same with and without exclusion for any given M . In contrast, the superiority-seeking motive predicts a clear directional difference.

4.1.1 Method

We recruited participants ($N = 118$) from a university-wide pool to take part in a decision-making study.²¹ Sessions were conducted with groups of $M \in \{4, 6, 8\}$, where each participant was assigned to an individual lab station numbered between 1 and 12. Participants earned \$15 as part of an unrelated study and were told that they may

²¹Instructions for all experiments can be found in the Online Appendix.

have the opportunity to purchase a unique good.²² The good was a custom T-shirt created specifically for experiments in our lab and was shown to participants across all sessions.²³

To measure valuations, we elicited participants’ maximum WTP for the good using incentive compatible mechanism described in Proposition 5, which corresponds to the classic BDM method. Here, after writing down their WTP (bid), the experimenter randomly drew a number between 1 and 15. This number served as the common ‘price’ P . If P was less than a participant’s WTP, then she would pay P to the experimenter and receive the object; if P was greater than the participant’s WTP, she would not pay anything and not get the object.

Participants were randomized into one of two conditions. In the Random Exclusion condition, the experimenter announced that he would roll a 12-sided die. If the outcome of the die roll matched a lab station number where a participant was sitting, that participant was not given the opportunity to bid for the good. The experimenter would roll the die until $K = \frac{M}{2} - 1$ participants were excluded and relinquished their bid sheets, where M was the group size ($K = \{1, 2, 3\}$ for $M = \{4, 6, 8\}$ respectively). The rest of the participants would write down their bids, which were collected by the experimenter. In the Baseline treatment, everyone in a session was given the opportunity to submit their WTP and potentially purchase the object.

At the end of the experimental session, those whose WTP exceeded P were paid $\$(15 - P)$ and received the object. Others were paid \$15.

4.1.2 Results

Figure 1 presents the means (Figure 1a) and distributions (Figure 1b) of participants’ WTP for the good by treatment. The distribution of WTP in the Random Exclusion treatment is shifted to the right of the distribution in the Baseline treatment. Consistent with this, the median WTP in the Random Exclusion treatment is \$5, double the median WTP in the Baseline treatment (\$2.50). Regressing WTP on a treatment dummy reveals a significant effect of exclusion ($\beta = 1.69, p = .03$).²⁴ These results provide direct support for superiority-seeking: consistent with Proposition 5, people significantly increase their willingness to pay for a good if they know that there are others whose intrinsic tastes for the good may well be in excess of theirs, but that such

²²The study was a simple effort task for which each participant earned \$15 for moving sliders across the computer screen, as in Gill and Prowse (2012).

²³We chose to use the custom T-shirt for two reasons. First, because it was created specifically for experiments in our lab, there was no salient anchor value. Second, participants who did not have the opportunity to bid for the shirt could not (easily) obtain it outside of the experiment. The second component was important for our exclusion manipulation.

²⁴Online Appendix Section 1.3.1 presents these results for each M separately. The effects are generally consistent, though some coefficients lose significance due to power issues.

intrinsic tastes are not met.²⁵




Figure 1: Willingness to Pay by Treatment

4.2 Study 2: Auctions

We now proceed to test the predictions of our framework in the competitive setting of first-price auctions. Our model makes predictions on the effects of exclusion that run counter to the classic results on competing for a private good. This facilitates a test of superiority-seeking in a setting that is both conceptually interesting and important for applications. Additionally, because the environment always involves the allocation of a single good, a positive effect of exclusion also allows us to rule out a host of alternative explanations related to changes in supply, such as scarcity being a signal of quality, as well as explanations based on consumption externalities or on scarcity of the opportunity to participate *per se*.

4.2.1 Method

Participants ($N = 274$) were recruited from a university-wide pool to take part in experiments on decision-making. As in the first study, sessions comprised groups of $M \in \{4, 6, 8\}$ and participants were assigned lab stations 1-12. Participants earned \$15 as part of the same unrelated study and then told that they may get the opportunity to participate in an auction. Conditional on having the opportunity, participants could

²⁵In Online Appendix Section 1.2, we use the data from Study 1 to structurally estimate the α parameter in this setting. Different methods provide consistent estimates of an α around 0.9. That being said, we do not want to put too much weight on these quantitative estimates as the exercise hinges on rational expectations about the distribution of valuations which people may exaggerate. The large magnitude may also be influenced by the public nature of exclusion in Study 1 (unlike the studies in Section 5).

use up to \$15 to bid on a good through a first price, sealed-bid auction. The good was a custom T-shirt as in Study 1. Participants would write down their bid privately on a sheet of paper. The highest bidder would receive the T-shirt and pay their bid. Everyone else would not receive the T-shirt and not pay anything.

Participants took part in one of three treatments. The Random Exclusion treatment was conducted in a similar way as in Study 1. Participants whose lab stations matched the outcome of the die roll would not have the opportunity to bid for the good and relinquished their bid sheets. The rest of the group reported their bids, which the experimenter then collected. The number of active bidders per session thus corresponded to $N = M - K$.

In the Non-Random Exclusion treatment, participants arriving to the lab were told about the T-shirt as in the other conditions but *not* about the auction. Each was then asked to indicate the extent to which she would want to own the good on a scale of 1-10. Once these scores were collected, participants learned about the auction and that $K = \frac{M}{2} - 1$ people from the group would not have the opportunity to participate. However, unlike in the Random Exclusion treatment, exclusion was based on participant's ex-ante liking of the good: the K individuals who least wanted to own the T-shirt were not given the opportunity to bid for it and this was made common knowledge. The actual liking scores of the participants were never revealed, only that those with the lowest were excluded. The number of active bidders here was thus exactly the same as in the Random Exclusion treatment, $N = M - K$.

We ran two versions of the Baseline treatment ($K = 0$). In both versions, the experimenter announced that everyone in this session would have the opportunity to bid on the shirt, i.e., $K = 0$. Participants then wrote down their bids, which were collected by the experimenter. The number of active bidders in this treatment was equal to the group size, $N = M$. The only difference between the two versions of the Baseline treatment was whether or not participants first indicated the extent to which they would want to own the good, matching the initial procedures of either the Random or Non-Random Exclusion treatments. This allowed us to test whether reporting this measure affected bids orthogonally to the treatment variation.

It is important to stress that the nature of exclusion—whether it was random or depended on intrinsic taste—was emphasized and made common knowledge as part of the experiment. Additionally, both the group size M and the number of excluded participants K were always emphasized in both written and verbal instructions. Care was also taken to make sure that participants left the lab one at a time and were away from the facility before the next participant departed. At the end of the experimental session, the highest bidder was paid \$15 minus her bid and received the shirt. All others were paid \$15.

4.2.2 Results

The average bid size was \$1.41 ($SE = 0.12$). There were no significant differences in bids between the two versions of the Baseline treatment ($p > .4$), indicating that reporting one’s ex-ante taste for the good did not meaningfully affect behavior. We thus pool data from the two versions for the analysis that follows.²⁶

We begin by examining the simple comparison of average bids in the Random Exclusion and Baseline treatments. A OLS regression with standard errors clustered at the session level reveals a sizable effect of exclusion (Random Exclusion=1) on bids ($\beta = 0.78, p < .01$).

Table 1 presents results from regressions that further test our predictions. We first regress bids on dummy variables corresponding to the number of active bidders N in the Baseline treatment. As shown in Column 1, the coefficients on both N dummies are positive. This provides evidence for the standard prediction outlined in Proposition 2 that, in the sale of private goods, increased competition leads to more aggressive bidding and higher prices.

In Column 2, we regress bids on dummy variables corresponding to the number of excluded K in the Random Exclusion treatment. The exclusion effect is economically significant: for example, the effect of excluding three out of eight people from the auction corresponds to 0.65 standard deviations of the mean bids in the study. Furthermore, the effect of excluding two people out of six from the auction ($M = 6, K = 0$ versus $M = 6, K = 2$) is roughly *double* the impact of adding two people in the Baseline treatment ($M = 4, K = 0$ versus $M = 6, K = 0$), i.e., the competition channel.²⁷ These results are consistent with the predictions of Proposition 2.

As outlined in Section 3, the superiority-seeking channel is muted in the Non-Random Exclusion treatment. When exclusion targets those with the lowest liking scores, in equilibrium the winner still has the highest realized private value in the group. At the same time, the classic effect, holding N constant, now predicts somewhat higher bidding due to positive selection on consumption utility. Comparing the relative effects of the Random versus Non-Random Exclusion treatments on average bids to that in the Baseline is thus a conservative test of the superiority-seeking motive. If $\alpha = 0$, then bids should be lower in Random Exclusion than in the Non-Random Exclusion treatment. A positive α alone does not imply a greater effect in the former than in the latter—only a sufficiently high α may.

Column 3 of Table 1 presents results from the Non-Random Exclusion treatment. The impact of exclusion here is much more muted, with smaller coefficients on K than in the Random Exclusion treatment. Column 4 compares the relative impact of the

²⁶ Similar results obtain when running the analyses separately for each version of the Baseline treatment.

²⁷ Online Appendix Section 1.3.1 presents these regressions separately for each M .

	(1)	(2)	(3)	(4)
$N = 6$	0.29 (0.22)			
$N = 8$	0.95*** (0.19)			
$K = 1$		0.67 (0.50)	-0.01 (0.48)	-0.08 (0.48)
$K = 2$		0.73** (0.35)	-0.26** (0.12)	-0.33** (0.15)
$K = 3$		1.28*** (0.18)	0.88 (0.53)	0.81 (0.54)
Random				-0.18 (0.24)
Random*($K = 1$)				0.85 (0.71)
Random*($K = 2$)				1.17*** (0.41)
Random*($K = 3$)				0.58 (0.59)
Constant	0.77*** (0.17)	1.12*** (0.11)	1.12*** (0.11)	1.19** (0.14)
N	142	210	206	274

*** : $p \leq 0.01$, ** : $p \leq 0.05$, * : $p \leq 0.1$. Standard errors clustered at the session level are reported in parentheses below each estimate. Column 1 reports the effect of group size N on bids in the Baseline treatments; Column 2 reports results comparing Random Exclusion to Baseline; Column 3 reports results comparing Non-Random Exclusion to Baseline; Column 4 compares the relative effects of Random versus Non-Random Exclusion.

Table 1: Effect of Group Size and Exclusion on Bids

two treatments to Baseline. The coefficients on the interactions between treatment and number excluded are all positive, suggesting that Random Exclusion has a larger impact on bids than Non-Random Exclusion. Since $\alpha = 0$ predicts a *smaller* effect in the former than in the latter, these results provide further evidence for superiority-seeking. Our third study, described in Section 5.1, provides more direct evidence for the role of beliefs about the desires of others in the superiority-seeking motive.

Above we focused on average bids between treatments. But we can also use our data to examine the predictions on expected revenue. To compute expected revenue, we ran a series of Monte Carlo simulations to generate bid distributions using the measured average bid and standard deviations for each group of active bidders N by

treatment. We draw N number of bids from these distributions to reproduce the type of data collected in the study, and take the maximum bid from each set of draws. This process is repeated 10,000 times for each combination of treatment and group size.

Expected revenues by treatment and group size M are presented in Figure 2. We find that the expected revenues from the Random Exclusion versus the Baseline treatment are 4.68 versus 3.84 for $M = 8$; 4.14 versus 3.03 for $M = 6$; and 4.04 versus 1.7 for $M = 4$. In line with Proposition 3, both in the Baseline treatment and in the Random Exclusion treatment, the seller's expected revenue is increasing in the number of active bidders. Exclusion, however, increases the seller's expected revenue by more than the increased competition effect of increasing the number of bidders with full inclusion.²⁸ Furthermore, in sharp contrast to the standard prediction absent superiority-seeking, the seller's expected revenue is considerably higher under random exclusion than under full inclusion both for each group size M and also when holding the number of active bidders N constant ($M = 6$ and $K = 2$ versus $M = 4$ and $K = 0$).

In contrast, expected revenues in the Non-Random Exclusion treatment are much closer to those in the Baseline treatment. For any fixed M and K these revenues are always below revenues in the Random Exclusion treatment, consistent with $\alpha > 0$. In line with Propositions 3 and 4, expected revenues in Non-Random Exclusion are roughly similar to those in the Baseline treatment (revenue equivalence follows for *any* α) and lower than under Random Exclusion for any fixed M and K (which follows for α sufficiently *high*; the opposite follows for $\alpha = 0$).




Figure 2: Expected Revenue by Group Size M Across Treatments

²⁸Starting at the same $M = 4$, expected revenue goes from 1.7 to 3.03 when N increases from 4 to 6 with full inclusion; expected revenue goes from 1.7 to 4.04 when N is decreased from 4 to 3 with exclusion.

4.3 Discussion

The preceding studies provide support for our model and rule out a number of alternative channels. First, there was always a single good being auctioned off across all treatments in Study 2, which assuages concerns about consumption externalities driving our results since these are held constant. Second, behavior in Study 2 also allows us to rule out mechanisms driven by variation in group size. Specifically, we can compare scenarios where the number of potential active bidders N is the same but there is presence versus absence of exclusion, i.e., $M = 6$ and $K = 2$ in Random Exclusion versus $M = 4$ and $K = 0$ in Baseline. The number of people excluded was also held constant across the Random and Non-Random Exclusion treatments. Since bids were higher in Random Exclusion than in both Baseline or Non-Random Exclusion, keeping the number of active bidders N constant, this allows us to also rule out differences in ‘scarcity of opportunity’ *per se* or social pressure as driving the observed effect.²⁹ Motives relating to joy of winning or signalling one’s income are either held constant across treatments or presumably imply more aggressive bidding in the Baseline than in the Random Exclusion treatment, predicting the opposite of the observed effect.

Another potential alternative is that our exclusion effect was driven by beliefs about the opportunities for resale. We can rule out a pure resale motive since the external resale opportunity of the winner is constant across treatments. One possibility is that the winner can resell only internally, that is, amongst the randomly excluded bidders. This seems unlikely to occur in practice because bidders were anonymous and care was taken that individuals left the lab one by one; moreover, we explicitly shut down the resale channel in the studies described in Section 5, which are conducted anonymously online. However, even if this unlikely scenario was the case in Study 2, given independent private values, random exclusion should not increase the seller’s revenue—in sharp contrast to what we find empirically. Given the standard regularity condition, whatever resale opportunities there may be, as long as $\alpha = 0$, it would still need to follow that the seller’s revenue satisfied $\Pi(M, K) < \Pi(M + 1, 0)$, e.g., [Klemperer and Bulow \(1996\)](#)—a prediction that is violated both in our setup and in the data. Instead, our results are consistent with the prediction of superiority-seeking, as described in Corollary 5, whereby randomly excluding bidders leads to a greater increase in expected revenue than expanding the number of bidders under full inclusion.

²⁹Notably, comparing behavior for constant N also allows us to rule out expectation-based reference dependence ([Kőszegi and Rabin, 2006](#)) as a driver of our effects. If equilibrium expectations are formed in the interim stage before bidding, then the only thing that matters is the number of active bidders N and thus bids should be same regardless of exclusion. If expectations are formed in the ex-ante stage for a given M and exclusion constitutes an out-of-equilibrium surprise relative to those expectations, then we believe that the spirit of the model would, if anything, predict the reverse of our findings. Importantly, expectation-based reference dependence predicts no treatment effects in the non-competitive settings of Studies 1, 3, and 4.

5 Robustness

The effects observed in the first two studies may still be impacted by aspects of the design, such as exclusion occurring face-to-face, exclusion being perceived as a signal of actual scarcity (Study 1), or other factors related to how exclusion was implemented. We designed our third and fourth studies to account for these potential confounds, to test the model’s predictions even more directly, and to further explore the economic significance of the superiority-seeking motive.

First, the new studies are conducted online in a fully anonymous setting. Any exclusion is thus completely anonymous, which should rule out image concerns that may have been present in the lab. Second, we employ a within-subject design to elicit valuations under different scenarios. Specifically, participants reported their valuations for the *same good* under scenarios that differed only in the number of others who would later be excluded from the opportunity to obtain the good. Since all exclusion scenarios were equally salient at the time that valuations were elicited, this allowed us to rule out concerns regarding differential attention as a driver of value (Krajbich et al., 2010; Li and Camerer, 2022; Towal et al., 2013). Third, we set up an online store exclusively for the study in order to credibly deliver unique goods to the participants remotely. Participants bid on a unique good—a print of a painting made by one of the authors. This setting highlights both the unique aspect of the product and alleviates concerns of exclusion being some function of actual scarcity.

5.1 Study 3: Exclusion on the Intensive Margin

The first two studies compared no exclusion to some exclusion as the main source of exogenous variation. This may have potentially introduced confounds, such as exclusion being a signal of scarcity, the feeling of ‘luck’ at not being excluded, or differences in group size at the time that valuations were elicited. The third study circumvents these issues by comparing valuations as the level of exclusion increases, conditional on some baseline level of exclusion. Participants reported their valuations for the same good under different levels of potential exclusion; decisions were incentivized by telling participants that one scenario would be selected at random and that decision would be played out ‘for real.’ This design choice makes transparent that the degree of exclusion is independent from scarcity or other participants’ valuations. Valuations are elicited *before* the actual scenario was realized, which precludes ex-post emotional factors such as ‘luck’ and the ‘joy of facing lower competition’ from impacting bids. Moreover, because the group size M was held constant at the time when valuations were elicited, this should further assuage concerns regarding differential social pressure as a driver of our results. The design also allows us to explore the superiority-seeking motive on

the intensive margin. Finally, we introduce treatment variation in information about whether the desire of those potentially excluded is higher or lower than one’s own in order to directly test the proposition that superiority-seeking increases with the excess desires of others.

5.1.1 Method

A group of participants ($N = 446$) were recruited from the Academic Prolific crowdsourcing platform to take part in experiments on decision-making.^{30,31} We refer to this group as ‘Active’ participants, which is in contrast to the ‘Passive’ participants whose role will be described below. Each Active participant was paid a base fee of \$1.00 for completing the study and was told that 1 in 10 would be randomly chosen to have their decisions played out for real.³²

Participants were first presented with an exclusive, unique good—a print of a painting made by one of the authors—and asked the extent to which they desired to own this good using the same method as in Study 2. Each was then told that they would have the opportunity to potentially purchase the good from an endowment of \$10 by reporting their maximum WTP, which was incentivized using the same common-price BDM mechanism as in Study 1. Further, these Active participants would be matched with three other Passive participants in the study.³³

Active participants were told that the Passive participants may have the opportunity to acquire the good and similarly report both their desire and their maximum WTP for it. Additionally, each of the Passive participants would have a separate coin assigned to each of them. The Active participants were then asked to report their WTP under three scenarios, with each scenario being equally likely to be selected and played out for real:

- **Scenario 1:** One of the coins would be flipped by the computer. If the coin flip landed on Tails, the assigned Passive participant would have the opportunity to purchase the good, depending on her WTP; if the coin flip landed on Heads, the assigned participant would not have the opportunity to purchase the good regardless of her WTP.
- **Scenario 2:** Two of the coins would be flipped by the computer and a second assigned Passive participant would be similarly affected by the outcome.

³⁰Experimental economics has increasingly used online crowdsourcing platforms. See, for example, Frydman and Jin (2022) and DellaVigna and Pope (2018).

³¹Pre-registration materials for this study can be found here <https://aspredicted.org/blind.php?x=CRD.N4B>

³²Paying a random subset of participants is a common practice in experimental economics and has generally yielded similar results as paying everyone (Charness et al., 2016).

³³The actual instructions used neutral language, with each participant being assigned a number.

- **Scenario 3:** All three coins would be flipped by the computer and the third assigned Passive participant would be similarly affected by the outcome.

Additionally, Active participants reported their WTP in one of three randomly assigned treatments. In the High Information treatment, each was told that the matched Passive participants had a desire for the good that was as high or higher than their own; in the Low Information treatment, the desires of matched participants would be as low or lower than the Active participant's. We also included a No Information treatment for robustness, where Active participants were not given any information about the desires of others in the group.

This design has several key features. First, the total size of the group remains constant across all conditions at $M = 4$. Second, as the number of coin flips increases, the greater the potential number of others excluded from the opportunity to own the good, i.e., K .³⁴ This allows us to test the model's predictions on the intensive margin. Third, our Informational treatments allow us to directly compare valuations when participants know that others have greater desire than their own (High Information) or know that they do not (Low Information). Specifically, the model predicts that Active participants will monotonically increase their WTP for the good as the number of coin flips increases, but only in the High and No Information treatments; the number of coin flips should not affect valuations in the Low Information treatment.

If an Active participant was chosen to have their decisions played out for real, they would be matched with a group of Passive participants who were recruited separately for the study. The group's choices would be realized as stated in the instructions. Payments were delivered digitally through the Prolific system. If a participant had purchased the art print, they received a code to be redeemed at an online store set up exclusively for the study. The online store was run through a third party that shipped the goods to participants upon redemption.

5.1.2 Results

Our analyses will focus on the High and Low Information treatments (where the relevant beliefs are directly controlled for) as these comprise the cleanest test of the framework, which predicts a boost in valuations as a function of beliefs about the excess valuations of others.³⁵ Figure 3 illustrates the results and Table 3 reports analyses from an OLS regression with standard errors clustered at the individual level.

³⁴In the one-coin scenario, the maximum is $K = 1$ and the mean is $K = 0.5$; in the two-coin scenario, the maximum is $K = 2$ the mean is $K = 1$; in the three-coin scenario, the maximum is $K = 3$, the mean is $K = 1.5$.

³⁵The results from the No Information treatment were similar to those in the High Information treatment. Consistent with the prediction of the model, all of the comparisons were significant but a bit less pronounced (see online Appendix 1.3).

In Column 1, we regress participants' bids on the number of coin flips in the High Information treatment. Increasing the number of those potentially excluded increases bids: compared to flipping only one coin (Scenario 1), flipping two coins increases bids by 32 cents ($p < .01$), and flipping three coins increases them by 68 cents ($p < .01$). This represents 0.12 and 0.25 standard deviations of the average bids in the study, respectively.³⁶

Column 2 presents the same results from the Low Information treatment. In contrast to the High Information treatment, exclusion has no significant effect when people think that the excluded have lower valuations than their own.

Column 3 presents results comparing the High and Low information treatments directly. Interacting both the two-coin and three-coin scenarios with the information treatment generates significant and sizable coefficients, showing that excess desire is necessary for superiority-seeking to emerge. Note that this comparison also allows us to rule out that the increase in valuations in the High Information treatment was driven by some artifact of the within-subject design; the only difference between the High and Low information conditions was whether or not the excluded had excess desire vis-à-vis the Active participant.

We chose to pair multiple Passive participants with each Active participant in order to explore the superiority-seeking motive on the intensive margin, holding the group size constant. This design also allowed us to examine both the existence and strength of superiority-seeking between-subjects. By looking at whether an individual increased her bids with the number of potentially excluded, and measuring the extent of this increase, we can capture both how many participants exhibited exclusionary preferences and the size of this effect on an individual level. To do so, we classify people who increased (decreased) their bids as the number of coin flips increased as exhibiting exclusionary (inclusionary) preferences. We classify those who did not change their WTP with the number of coin flips as neutral.³⁷

³⁶To put these effect sizes in perspective, we benchmarked the effect of exclusion against a factor that has been widely considered to be economically important in driving valuations—the endowment effect. We ran a version of the classic endowment effect paradigm from [Kahneman et al. \(1990\)](#) adapted to our setting ($N = 183$), using the same art print as in Study 3. We found that being endowed with the unique art print led to a 70 cent increase in valuations relative to not being endowed with it ($p < .05$). This suggests that in our setting, the classic endowment effect is similar in magnitude to increasing the number of potentially excluded from one to three.

³⁷Our classification compared bids in Scenario 3 to Scenario 1. If bids in the former were strictly greater (smaller) than in the latter, the participant was classified as exhibiting exclusionary (inclusionary) preferences. In our sample, 98.9% of bids were weakly monotonic with respect to this classification, e.g., someone classified as having exclusionary preferences did not bid strictly lower in Scenario 2 than Scenario 1. We drop participants who reported a WTP of \$0 or \$10 in Scenario 1 from the classification analysis as their valuations are subject to censoring concerns (i.e., those reporting \$0 (\$10) cannot decrease (increase) their WTP further in Scenario 3). Results are qualitatively the same without this restriction (e.g., exclusionary preferences are still by far the biggest category in the High Information treatment).

	(1)	(2)	(3)
Two-Coin (=1)	0.32*** (0.09)	-0.13 (0.09)	-0.13 (0.09)
Three-Coin (=1)	0.68*** (0.16)	-0.09 (0.14)	-0.09 (0.14)
High Information (=1)			0.35 (0.30)
Two-Coin*High Information			0.46*** (0.12)
Three-Coin*High Information			0.78*** (0.21)
Constant	3.67*** (0.22)	3.33*** (0.21)	3.33*** (0.21)
<i>N</i>	444	453	897

*** : $p \leq 0.01$, ** : $p \leq 0.05$, * : $p \leq 0.1$. Standard errors clustered at the individual level are reported in parentheses below each estimate. Column 1 reports the relationship between the number of coin flips and WTP in the High Information treatment. Column 2 reports the relationship between the number of coin flips and WTP in the Low Information treatment. Column 3 compares the High and Low Information treatments.

Table 2: Effect of Exclusion on WTP




Figure 3: Average Bids by Information and Level of Exclusion

In the High Information treatment, 51% of people can be classified as exhibiting strict exclusionary preferences. This is the largest group, with 29% exhibiting neutral preferences and only 20% exhibit the opposite tendency.³⁸ Importantly, these proportions do not reflect noise or artifacts of the experimental design: consistent with the theory, the proportion of people with exclusionary preferences is substantially lower—nearly half the size—in the Low Information treatment (26%; $p < .01$).

This classification also allows us to examine the strength of the superiority-seeking motive. To do so, we replicate Column 1 of Table 3 conditional on exhibiting exclusionary preferences. Active participants increased their WTP for the good by \$1.10 when going from one to two coins, and by \$2.44 when going from one to three. These are sizable effects, representing .40 and .89 standard deviation, respectively, from the mean WTP in the study.

5.2 Study 4: Shifting Demand Through Exclusion

We designed our last study to mirror the setting of Proposition 1 via a within-subject design and explore whether, in line with our prediction, a seller may increase profits through random exclusion compared to the classic optimal mechanism of price-based exclusion which targets only those with the lowest valuations.

5.2.1 Method

The paradigm builds on the third study with several notable differences. Participants ($N = 100$) were recruited from the Academic Prolific crowdsourcing platform to take part in experiments on decision-making.³⁹ Every participant was paid a base fee of \$1.00 for completion. Each also had a \$10 endowment to use for decisions in the study, and all participants had their decisions carried out for real.

Participants were informed that 100 individuals had been recruited for this study. Each then reported her WTP for the good—incentivized in the same way as in Study 3—in two (within participant) scenarios. In the No Exclusion scenario all participants who finished the study would be able to purchase the good depending on their WTP. In the Exclusion scenario 60% of participants would be randomly selected to have the opportunity to purchase the good; the rest would not have this opportunity. A computer would then flip a coin to determine which scenario would be played out for

³⁸Notably, the proportion exhibiting exclusionary preferences is comparable to the proportion of people exhibiting loss aversion and the endowment effect—two of the most well-studied phenomena in behavioral economics. Studies that explore heterogeneity in loss attitudes find that 53% (Dean and Ortoleva, 2019) of participants are loss averse, while Goette et al. (2019) found that 57% of people exhibit the endowment effect.

³⁹Pre-registration materials for this study can be found here https://aspredicted.org/blind.php?x=D3B_X6X

real, with Heads corresponding to Exclusion and Tails corresponding to No Exclusion.⁴⁰ Note that, like in Study 3, the within-subject design makes it transparent that exclusion is not a function of actual scarcity.

5.2.2 Results

We first look at whether the prospect of exclusion affected average WTP for the good. Regressing WTP on a treatment dummy (Exclusion=1) with clustered standard errors reveals a significant effect: participants increased their WTP for the good by nearly a dollar ($\beta = 0.91; p < .01$) for the Exclusion case relative to the No Exclusion case. This represents .27 standard deviations from the mean WTP.

We can then calculate the demand curves facing a monopolist both with and without exclusion. Below, we report the number of participants who would be willing to purchase the good at each price for both scenarios. The two demand curves are presented in Figure 4 below.




Figure 4: Demand as a Function of Exclusion

The results support our theoretical predictions, Propositions 1 and 5, and lend credence to revenue-generating non-price-based exclusion. Demand shifted to the right as a function of exclusion.⁴¹ Participants' median WTP increased by nearly 50%—from \$4 to \$6—as a function of others being excluded from the opportunity to purchase the

⁴⁰Unlike in Study 3 where all Active participants had the opportunity to purchase the good, if the outcome of the coin flip was Heads, only participants who were amongst the 60% selected would have this opportunity.

⁴¹Kolmogorov–Smirnov test for equality of distributions confirms that this shift is significant ($p = .02$).

good. Furthermore, as long as the marginal cost of supplying one good is \$2 or more—which it was in our setting—the seller’s optimal profit under exclusion is *higher* than under no exclusion. This, despite getting rid of 40% of potential demand and thus having a non-selectively 40% smaller customer base.

We can also perform a similar classification exercise as in Study 3, tagging those whose bids in the Exclusion scenario are strictly higher (lower) than in the No Exclusion scenario as having exclusionary (inclusionary) preferences. Those with no differences in bids are classified as neutral. We observe a similar breakdown as Study 3, with 53% exhibiting exclusionary, 25% exhibiting neutral, and 22% exhibiting inclusionary preferences. Looking at the size of the superiority-seeking motive, we find that conditional on exhibiting exclusionary preferences, people’s median WTP for the good increases by nearly 80%—from \$5 to \$9—when others are barred from acquiring the good. We believe that this striking empirical demonstration provides evidence for the potential effectiveness of exclusion as a rent generating tool for firms.

6 Discussion

6.1 Related Literature

A rich literature emphasizes the presence of social motives in people’s preferences. A first generation of social preference models focuses on direct consumption externalities, e.g., (Becker, 1974). The next wave of models attempted to explain empirical evidence on costly punishment behavior, by considering preferences over *relative* allocations. Models of inequity aversion (Fehr and Schmidt, 1999) and competitive preferences (Charness and Rabin, 2002) over money assume that people dislike unequal allocations or actively prefer allocations that put them ahead, respectively. A third wave of theory considers whether prosocial or antisocial choices may be a function of imperfect information and signalling rather than underlying preferences, e.g., Bénabou and Tirole (2006, 2011).

In these models, utility is defined over the consumption, money, or the beliefs of others. This is in contrast to our approach which defines utility over others’ unmet desires. This distinction is important because it generates novel predictions on the effects of exclusion. For example, in the auction setting, only the winner receives the item while others do not. As a result, there are no changes in consumption externalities and no differences in the relevant informational asymmetries, as the allocation mechanism is common knowledge. Hence, these alternative models will not predict our findings, while superiority-seeking does.

In a line of work closer to our own, Frank (1985) models ‘keeping up with the Joneses’ effects through a demand for positional, or status, goods. The existence of such

goods is taken as a primitive and people not only care about their personal consumption utility from positional goods, but also the hierarchy of *observable* consumption amongst others. In particular, people’s utility from a positional good is a function of the percentile ranking of their own consumption of the good relative to the overall population’s consumption of it. If people make choices non-cooperatively, this may then lead to overconsumption of positional goods compared to non-positional goods. ⁴²

Our framework is conceptually distinct from this literature: in contrast to a direct consumption externality, the superiority-seeking motive is defined over the unmet desires of others for what one has rather than over the consumption space independent of those desires. Importantly, our mechanism makes distinct empirical predictions. As noted above, consumption externalities such as Frank (1985) would predict a null effect in Study 2. At the same time, positional externalities are consistent with the implications of superiority-seeking, which imply an incentive to surpass the consumption of others, in quality or kind, if this allows a person to increase the unmet desires of others for what she has. Our framework may thus be viewed as a psychological foundation for the demand for positional or status goods, while making a host of additional predictions for individual decision-making and market behavior. We expand on this further below.

6.2 Broader Implications

In this section, we explore some economic consequences implied by our framework for markets and political economy.

Exclusive Access and Artificial Scarcity. As captured by Proposition 1 and Corollary 4, firms and marketers can extract greater rents by artificially rationing the actual or perceived supply of goods and services even while employing a uniform price. Examples of such practices abound. Established venues face persistent excess demand at the going price in lieu of expanding capacity or raising the price despite ample opportunities to do so. Nightclubs grant random access to some but not others despite there being enough room for all.⁴³ Firms regularly introduce new consumer products at artificially low quantities despite overwhelming demand, e.g., Gmail and Facebook rationing invitations (Wortham, 2011). As another example, the major Chinese cell-

⁴²See also Heffetz and Frank (2011) for a review of the literature on status-seeking, with status goods as a special case.

⁴³A classic example was the nightclub Studio 54 in Midtown Manhattan, described by Andy Warhol as “dictatorship on the door but a democracy on the dance floor.” Here, “guestlist and celebrities were allowed in a separate entrance through the back door but the general public lined the streets in the hope to be given the chance to pay to enter the club. ‘When Steve was running the club he would be very difficult with people on the front door,’ says Mark Fleischman, previous Studio 54 owner and author of the recently released Inside Studio 54 book. ‘He’d say, “you can’t come in, I don’t like your shirt, get out.” People would wait for hours and not be let in.’” (Wray, 1989).

phone manufacturer OnePlus distributed their flagship OnePlus One phone by first soliciting a large list of people who were interested in buying it and then randomly allocating the opportunity to a select few. As discussed in Section 1, marketers work hard to emphasize such ‘artificial’ scarcity when advertising products. Techniques include limited time offers, highlighting that there are ‘only’ a few goods left (common in online marketplaces), or that a price discount is offered to a select few and only allowing people with last names beginning with certain letters of the alphabet to purchase a good (a classic infomercial practice). Practitioner guides note that such scarcity techniques are effective in increasing demand by making the eventual owners ‘feel powerful’ over those who want the same product but cannot have it.⁴⁴

As discussed in Becker (1991), these phenomena are puzzling from the perspective of standard economics since a firm should either raise the price and/or expand capacity; firms should not choose to restrict supply while profit opportunities are seemingly positive. His paper highlights the widespread prevalence of such artificial scarcity but ultimately cannot explain it. In contrast, our model rationalizes the use of artificial scarcity and excess demand at the posted price as a way to increase profits.

Price-based Exclusion and Veblen Goods. Economists have long considered the prevalence of status goods and so-called ‘Veblen effects,’ where the demand for certain goods increases as its price increases. The goods that display this apparent violation of the law of demand are often characterized by their exclusivity and luxury. A line of work relates ‘Veblen effects’ to the propensity to signal one’s income status via observable consumption for downstream non-pecuniary benefits. For example, Bagwell and Bernheim (1996) consider a setting where people allocate their income across two types of goods: observable ‘conspicuous’ goods, such as a sport car, and less observable non-conspicuous, goods, such as private home decor or art displayed at home. A ‘Veblen effect’ is said to arise when rich people consume a higher-priced but otherwise equivalent version of the ‘conspicuous’ good. However, the conditions which generate such effects also generate fairly counter-intuitive predictions. For example, a more efficient method of signaling income would be for richer people to destroy resources publicly, or simply just post their incomes. The framework suggests that firms have an incentive to produce *perfect* substitutes that only differ in price, creating a ‘budget’ product and a ‘luxury’ product of the same quality. Finally, the purely signaling explanation for ‘Veblen effects’ and luxury goods has a difficult time rationalizing the often private consumption of the associated goods and the persistence of high prices despite observationally similar alternatives, e.g., real versus fake diamonds.

⁴⁴For example, SUMO’s guide to scarcity marketing states that “Scarce items make people feel powerful: Snagging a scarce item means you have access to something other people want but can’t have—which gives the owner power” (<https://sumo.com/stories/scarcity-marketing>).

Our framework provides an alternative and perhaps complementary channel to signalling. In the presence of heterogeneous income, people’s demand for a luxury good may increase with its price irrespective of observability with regards to one’s own income status or consumption, as a high price creates exclusion. Increasing prices lead consumers who may very much desire the good to no longer be able to obtain it due to tighter liquidity constraints, which generates a superiority-seeking boost amongst those who may like the good less but can still afford it. This can lead to an increase in total demand. A simple stylized example provides intuition.

Example 3. *Let there be M consumers, each is rich with probability γ and poor otherwise. A monopolist produces a good at zero marginal cost. Each person has a unit demand and her consumption utility is an i.i.d. draw from $F(v)$ over $[0, \bar{v}]$. The only ex-ante difference between rich and poor is that poor consumers face a tighter budget constraint. For simplicity, suppose that the poor cannot spend more on the good than some number $Z < \bar{v}$.*

If the price is less than Z , total expected demand is $M[1 - F(p)]$. If it is greater, there is no demand by the poor, but the expected demand is now $\gamma M[1 - F(p - \alpha g(p))]$ for some function $g(p) > 0$ as long as $p < \bar{v}$ reflecting the utility from superiority-seeking; rich consumers with an intrinsic taste well below p will now also want to buy the product. It is straightforward to see that, depending on F and γ , demand may indeed be higher than when the price is below Z , emblematic of the standard ‘Veblen effect.’⁴⁵

The above stylized example also illustrates why, in contrast to the prediction of the signaling account, a firm may not want to engage in third-degree price discrimination—even if the company could sell to the verifiably rich and poor at different prices and it was clear at which price a particular good was bought. Instead, the firm would be better off shutting down the market for poorer individuals. Producing low-priced ‘budget’ substitutes introduces the prospect of low-income individuals having access to the same goods as richer individuals, which would diminish the impact of superiority-seeking and crowd out demand in the latter group. This rationalizes why luxury firms rarely advertise budget products under the same umbrella; when a firm does offer both low and high priced goods, the former is typically advertised under a different brand

⁴⁵Note that in our basic setup we assumed, as standard in market and mechanism design settings, that utility was commonly quasi-linear in money. In the context of more binding budget constraints for the poor, we can still maintain such a quasi-linearity assumption for the poor and apply our framework if, as is introduced by [Weitzman \(1977\)](#) and developed in [Friedman and Sákovic \(2015\)](#), we allow the marginal utility of money to be higher for the poor than for the rich; e.g., they are known to face higher interest-rates (potentially infinite ones). Note also that the logic extends to the case where there is a more general distribution of budget constraints across the potential consumers. Then the induced demand may well be smoothly upward sloping in price over a larger and continuous range of prices.

and stresses accessibility, while the latter stresses quality (e.g., Armani Exchange versus Giorgio Armani).

Importantly, our model also makes distinct predictions on the *types* of objects that are more likely to display ‘Veblen effects’. Unlike in the signaling account, intrinsic quality now plays an important role—a Ferrari is more likely to be a Veblen good than a Nissan Altima because it is a better quality car. A higher quality good is associated with a greater intrinsic taste for it and will be linked to higher levels of superiority boost when this desire is unmet. This is in contrast to pure signaling motives where perceptions of quality do not factor into ‘Veblen effects’ *per se*. Additionally, our framework allows for exclusive goods to be enjoyed in private, such as intimate dinners at an expensive restaurant, art collections housed in a private gallery, or luxury amenities of a gated community. Perceptions of exclusivity rather than observability of one’s own consumption is the key driver for increased demand in our framework.

We do not claim that the observability of consumption is not an important factor in the demand for luxury goods. Work by, for example, [Heffetz \(2011\)](#) and [Bursztyn et al. \(2018\)](#), have shown that visibility plays a significant role in the consumption of some premium goods. In fact, visibility may amplify the utility boost in our model by ramping up desire amongst those who cannot afford the good or by shaping the social context. We also do not aim to minimize signaling as a potential motive. Rather, we argue that the utility boost from unmet excess desires provides a distinct and potentially complementary channel which further increases the predictive power for which goods are more or less likely to generate such ‘Veblen effects’ or become status goods.

Redistribution. Redistributive policies such as progressive taxation are common tools for mitigating income inequality. However, political scientists and economists are often puzzled by opposition to these policies by the U.S. electorate, particularly amongst the poor and lower middle-class.⁴⁶ Prior work has argued that such opposition may stem from the ‘prospect of upward mobility’ (e.g., [Benabou and Ok 2001](#)), whereby people would prefer to avoid higher taxes given their prospects of moving up higher on the income ladder, or motivated beliefs about the potential of upward mobility which function to counteract limited willpower ([Bénabou and Tirole, 2006](#)). In the presence of income heterogeneity, superiority-seeking offers a potential complementary and distinct motive for the opposition to some re-distributive policies—even by those who may benefit from it, e.g., amongst the poor-but-not-poorest individuals of a community. These policies may decrease the overall utility associated with the respective group’s consumption despite providing greater social insurance.

Consider the prospect of increasing the minimum wage. A higher minimum wage

⁴⁶For recent evidence, see, e.g., [Kuziemko et al. \(2015\)](#).

may have positive implications for those currently earning below the minimum wage; they can purchase goods that previously only those who are richer could afford. Our framework predicts why opposition may be strongest amongst those who earn just one bracket above the current minimum wage. To see the intuition, consider the group earning just one income notch above the current minimum wage. Similar to the Veblen logic described above, this group is currently deriving a utility boost from consuming products that the lower income group—who earn at the current minimum wage—cannot afford. Increasing the minimum wage to the level of their income would eliminate this boost, which leads the poor-but-not-poorest earning group to oppose the policy. In fact, the decrease in the superiority boost may well be the largest for those who are currently just one income notch from the bottom. Given heterogeneity in tastes, others closer to the top of the income distribution will still continue to enjoy various goods that those below them like more, but cannot afford.⁴⁷ This prediction is consistent with the findings of Kuziemko et al. (2014), who present survey evidence collected by a marketing group that people just one income notch above the minimum wage threshold oppose a minimum wage increase the most even after controlling for demographics.

Immigration, Nationalism, and Barriers to Trade. The mechanism described in the paper is also broadly consistent with the phenomenon whereby ‘natives’ may want to limit the rights of ‘immigrants’, even if such exclusion comes at a material cost to the former. By restricting access to certain rights and institutions that some immigrants may desire more, natives can derive extra utility through the superiority-seeking motive. This also helps explain the familiar notion of ‘pulling up the ladder,’ whereby people who have recently immigrated oppose further immigration.⁴⁸ Even if further immigration were to increase their own material well-being and productive social networks, recent immigrants who know others who would like to immigrate as well—thereby gaining access to opportunities and institutions available in the host country—may oppose additional immigration because this would lead to a drop in their own utility (which is derived from distinctive access to this set of benefits).

Politicians appear aware of these fears, describing access to national institutions through the perspective and desires of those outside of the nation. Often this rhetoric takes the form of highlighting outsider desire, such as in the United Kingdom where the National Health Service is often referred to as the ‘envy of the world,’ or in the

⁴⁷Specifically, even holding the social context constant, the superiority-seeking boost derived by those closer to the top of the distribution is more likely to come vis-à-vis people who are not at the bottom of the income ladder, and hence their utility will not be directly affected by redistributive policies targeting the bottom of the distribution.

⁴⁸For example, a recent Pew Research survey showed that half of all foreign-born whites stated that new immigrants threaten US values rather than strengthen them. A quote from a recent immigrant, who came in 2003, sums up this sentiment “I think that enough immigrants entered this country” (<https://archive.sitrib.com/article.php?id=4132971&itype=CMSID>).

United States, where the US economy is celebrated as 'being the envy of the entire world.'⁴⁹

Although our analysis has focused on the individual's superiority-seeking vis-à-vis her social context, it can also apply to the joy of identifying with a group and the rivalry between groups. In particular, people may enjoy identifying with a group they can belong to and derive pleasures from the goods and attributes this group as a whole possesses. The value of in-group identification is then amplified if that group possesses attributes or consumes goods that members of another group would like more, but do not have. Indeed, given the superiority-seeking motive, it is the exclusion of out-of-group members from such goods or attributes which causes a form of pride and generates a utility boost from one's group identity. To protect the value of such group identities, maintaining exclusion is important, which generates psychological barriers to inter-group trade.

Discrimination and Social Stratification. Similarly, the above logic can help rationalize aspects of social exclusion and 'taste-based' discrimination. In a framework termed *stratification economics*, members of social groups compete over relative positions in exogenous social hierarchies (Darity Jr et al., 2017, 2015). Higher positions provide members with a number of privileges including exclusive access to a broad category of *club goods*. Here, discrimination is a 'rational' response by dominant groups to maintain access to these privileges, serving as a tool for exclusion so that their own supply is protected. This is consistent with classic examples of club goods such as country clubs or exclusive residential communities being historically marked by discrimination based on socially-constructed groups.⁵⁰

Our framework provides a distinct and complementary account through the psychological motive of superiority-seeking. The model predicts exclusionary and discriminatory policies even in the absence of exogenous hierarchies or the club goods being rivalrous on the relevant margin. Specifically, majority groups may employ discriminatory policies in order to boost their own private utility from consuming private goods or possessing certain attributes. That is, the unmet desires of the excluded or persecuted minority yield additional utility benefits associated with consumption of those goods, which increases incentives for discrimination. Exclusion based on salient (or easily observable) characteristics such as race or ethnicity further facilitates superiority-seeking by members of the majority group. In this account, the preferences that lead to dis-

⁴⁹See, for example, Riley-Smith (2019). President Biden referred to democratic elections in similar terms: "But that patience has been rewarded now for more than 240 years, the system of governance has been the envy of the world." (Buncombe, 2020).

⁵⁰For example, see the real estate ad for properties in La Jolla, CA, a community marked by a 'gentleman's agreement' not to sell housing to Jews: "The very fact that you live in La Jolla puts you in a special class"(Stratthaus, 1996). Private golf clubs have a long history of race-based discrimination (Sawyer, 1992).

parate treatment do not arise from some innate disutility from interacting with others of a certain physical or ethnic characteristic; rather, they arise due to superiority-seeking, whereby external and observable aspects of others serve as a coordination device or as modes of ‘valuable’ identification as described before. The logic of how superiority-seeking reinforces social discrimination and stratification is similar to the other cases of non-price-based exclusion in the preceding discussion.

7 Conclusion

Our paper proposes that an individual’s valuation of consuming an item or possessing an attribute is boosted by a comparative term reflecting others’ unmet desire for it. Our model of imitative superiority-seeking helps explain a host of market anomalies and generates novel predictions for competition and political economy. We present experimental evidence that supports the predictions of the framework across several economically-important environments. Our model rationalizes the use of artificial restrictions to supply (non-price-based exclusion) by firms, scarcity marketing, and naturally generates ‘Veblen effects.’ The framework also provides a novel motive for attitudes against redistribution and immigration, and points to a distinct psychological mechanism for inter-group discrimination that has broad implications for social rivalry in a variety of key economic and political contexts.

Future research can greatly refine and expand the predictions of the motive introduced in our paper. From a theoretical perspective, preference for exclusion may have important implications for a variety of pricing and allocation problems. It may also be useful to consider factors that shape one’s social context and determine the domains in which superiority-seeking is most pronounced. Similarly, further studies could consider social institutions, such as systems of honour or moral exclusion that may both amplify and channel superiority-seeking from one domain to another. We believe that this motive may well be key for understanding a host of issues in political economy.

A natural question emerges regarding how other aspects of social behavior are related to the superiority-seeking motive. As an initial exploration, we re-ran the No Information treatment of Study 3 ($N = 150$) and included measures of other social behavior that have been studied in the past, such as altruism, cooperation, and prosocial and anti-social punishment.⁵¹ We classified people as having exclusionary, neutral, or inclusionary preferences based on the method used in Study 3. Superiority-seeking was not correlated with any of the ‘prosocial’ measures such as altruism, cooperation,

⁵¹Altruism was captured by giving in the dictator game (Forsythe et al., 1994), cooperation was captured by behavior in the prisoner’s dilemma and contributions in the public goods game (Camerer, 2011), and prosocial and anti-social punishment was captured by the strategy method in response to contributions in the public goods game (Herrmann et al., 2008). All decisions were incentivized by selecting one choice at random and having it play out for real.

or prosocial punishment. We did find a significant, albeit small, correlation between superiority-seeking and anti-social punishment.⁵² This suggests that preferences for exclusion are conceptually distinct from other social motives but may share some underlying psychology with anti-social behavior.

Our model has a number of limitations. An important ingredient of our model is one's social context. We do not endogenize this social context, which can be affected by attention, memory, and salience of the comparison set (see, for example, work by [Bordalo et al. \(2013\)](#), [Bordalo et al. \(2020\)](#), [Smith and Krajbich \(2019\)](#), and [Gabaix \(2019\)](#)). Furthermore, strategic agents, such as firms or politicians, may try to directly influence it by the narratives they promote and the rhetoric they engage in. Additionally, in many contexts superiority-seeking may be veiled by factors outside of our model, such as the desire to conform to norms and forms of behavior that would, for example, prevent people from explicitly paying for social exclusion. Future research should also explore the psychological factors that determine how superiority-seeking interacts with broader economic conditions.

⁵²The correlation between having exclusionary preferences and being willing to engage in anti-social punishment is 0.14 ($p < .01$).

References

- Saint Augustine. *The City of God*. Hendrickson Publishers, 2009.
- Laurie Simon Bagwell and B Douglas Bernheim. Veblen effects in a theory of conspicuous consumption. *The American Economic Review*, pages 349–373, 1996.
- Gary S Becker. A theory of social interactions. *Journal of Political Economy*, 82(6):1063–1093, 1974.
- Gary S Becker. A note on restaurant pricing and other examples of social influences on price. *Journal of Political Economy*, 99(5):1109–1116, 1991.
- Roland Benabou and Efe A Ok. Social mobility and the demand for redistribution: the poum hypothesis. *The Quarterly Journal of Economics*, 116(2):447–487, 2001.
- Roland Bénabou and Jean Tirole. Incentives and prosocial behavior. *The American Economic Review*, 96(5):1652–1678, 2006.
- Roland Bénabou and Jean Tirole. Identity, morals, and taboos: Beliefs as assets. *The Quarterly Journal of Economics*, 126(2):805–855, 2011.
- Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Salience and consumer choice. *Journal of Political Economy*, 121(5):803–843, 2013.
- Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Memory, attention, and choice. *The Quarterly Journal of Economics*, 135(3):1399–1442, 2020.
- Andrew Buncombe. ‘the vote is sacred’: Biden says he has no doubt he has won but democracy requires patience. *The Independent*, 2020. URL <https://www.independent.co.uk/news/world/americas/us-election-2020/biden-win-votes-count-2020-election-trump-b1626691.html>.
- Leonardo Bursztyn, Bruno Ferman, Stefano Fiorin, Martin Kanz, and Gautam Rao. Status goods: experimental evidence from platinum credit cards. *The Quarterly Journal of Economics*, 133(3):1561–1595, 2018.
- Modibo K Camara. Computationally tractable choice. Technical report, Tech. rep., Working Paper, 2021.
- Colin F Camerer. *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press, 2011.
- Gary Charness and Matthew Rabin. Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, 117(3):817–869, 2002.

- Gary Charness, Uri Gneezy, and Brianna Halladay. Experimental methods: Pay one or pay all. *Journal of Economic Behavior Organization*, 131:141–150, 2016. ISSN 0167-2681. doi: <https://doi.org/10.1016/j.jebo.2016.08.010>. URL <https://www.sciencedirect.com/science/article/pii/S0167268116301779>.
- Robert B Cialdini. *Influence: The psychology of persuasion*, volume 55. Collins New York, 2007.
- William A Darity Jr, Darrick Hamilton, and James B Stewart. A tour de force in understanding intergroup inequality: An introduction to stratification economics. *The Review of Black Political Economy*, 42(1-2):1–6, 2015.
- William A Darity Jr, Darrick Hamilton, Patrick L Mason, Gregory N Price, Alberto Davila, Marie T Mora, and Sue Stockly. Stratification economics: a general theory of intergroup inequality. *The Hidden Rules of Race: Barriers to an Inclusive Economy*, pages 35–51, 2017.
- Mark Dean and Pietro Ortoleva. The empirical relationship between nonstandard economic behaviors. *Proceedings of the National Academy of Sciences*, 116(33):16262–16267, 2019.
- Stefano DellaVigna and Devin Pope. What motivates effort? evidence and expert forecasts. *The Review of Economic Studies*, 85(2):1029–1069, 2018.
- Ernst Fehr and Klaus M Schmidt. A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3):817–868, 1999.
- Gregory W Fischer, Mary Frances Luce, and Jianmin Jia. Attribute conflict and preference uncertainty: Effects on judgment time and error. *Management Science*, 46(1):88–103, 2000.
- Robert Forsythe, Joel L Horowitz, Nathan E Savin, and Martin Sefton. Fairness in simple bargaining experiments. *Games and Economic Behavior*, 6(3):347–369, 1994.
- Robert H Frank. The demand for unobservable and other nonpositional goods. *The American Economic Review*, 75(1):101–116, 1985.
- Daniel Friedman and József Sákovics. Tractable consumer choice. *Theory and Decision*, 79(2):333–358, 2015.
- Cary Frydman and Lawrence J Jin. Efficient coding and risky choice. *The Quarterly Journal of Economics*, 137(1):161–213, 2022.
- Xavier Gabaix. Behavioral inattention. In *Handbook of Behavioral Economics: Applications and Foundations 1*, volume 2, pages 261–343. Elsevier, 2019.

- David Gill and Victoria Prowse. A structural analysis of disappointment aversion in a real effort competition. *The American Economic Review*, 102(1):469–503, 2012.
- René Girard. Deceit, desire and the novel: Self and other in literary structure, trans. yvonne freccero, 1965.
- René Girard. *Oedipus unbound: selected writings on rivalry and desire, edited by Mark R. Anspach*. Stanford University Press, 2004.
- Lorenz Goette, Thomas Graeber, Alexandre Kellogg, and Charles Sprenger. Heterogeneity of gain-loss attitudes and expectations-based reference points. *Available at SSRN 3589906*, 2019.
- Milton Harris and Artur Raviv. A theory of monopoly pricing schemes with demand uncertainty. *The American Economic Review*, 71(3):347–365, 1981.
- Ori Heffetz. A test of conspicuous consumption: Visibility and income elasticities. *Review of Economics and Statistics*, 93(4):1101–1117, 2011.
- Ori Heffetz and Robert H Frank. Preferences for status: Evidence and economic implications. In *Handbook of Social Economics*, volume 1, pages 69–91. Elsevier, 2011.
- Benedikt Herrmann, Christian Thoni, and Simon Gächter. Antisocial punishment across societies. *Science*, 319(5868):1362–1367, 2008.
- Thomas Hobbes. Man and citizen. originally published, 1642 and 1658, 1998.
- Daniel Kahneman, Jack L Knetsch, and Richard H Thaler. Experimental tests of the endowment effect and the coase theorem. *Journal of Political Economy*, 98(6):1325–1348, 1990.
- Paul Klemperer and J Bulow. Auctions versus negotiations. *The American Economic Review*, 86(1):180–94, 1996.
- Botond Köszegi and Matthew Rabin. A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4):1133–1165, 2006.
- Ian Krajbich, Carrie Armel, and Antonio Rangel. Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, 13(10):1292–1298, 2010.
- Ilyana Kuziemko, Ryan W Buell, Taly Reich, and Michael I Norton. “last-place aversion”: Evidence and redistributive implications. *The Quarterly Journal of Economics*, 129(1):105–149, 2014.

- Ilyana Kuziemko, Michael I Norton, Emmanuel Saez, and Stefanie Stantcheva. How elastic are preferences for redistribution? evidence from randomized survey experiments. *The American Economic Review*, 105(4):1478–1508, 2015.
- Chuan-Zhong Li and Leif Mattsson. Discrete choice under preference uncertainty: an improved structural model for contingent valuation. *Journal of Environmental Economics and Management*, 28(2):256–269, 1995.
- Xiaomin Li and Colin F Camerer. Predictable effects of visual salience in experimental decisions and games. *The Quarterly Journal of Economics*, 137(3):1849–1900, 2022.
- Chavie Lieber. Why fashion brands destroy billions’ worth of their own merchandise every year. *accessed on July, 17, 2018*. URL <https://www.vox.com/the-goods/2018/9/17/17852294/fashion-brands-burning-merchandise-burberry-nike-h-and-m>.
- Paul R Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica*, pages 1089–1122, 1982.
- Andrea Marie Moreno. *Hypebeast Fashion: Why Streetwear Is the New Luxury Fashion*. Fashion Institute of Technology, State University of New York, 2020.
- Diana C Mutz. Status threat, not economic hardship, explains the 2016 presidential vote. *Proceedings of the National Academy of Sciences*, 115(19):E4330–E4339, 2018.
- Roger B Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1): 58–73, 1981.
- Wolfgang Pesendorfer. Design innovation and fashion cycles. *The American Economic Review*, pages 771–792, 1995.
- Charles R Plott and Kathryn Zeiler. The willingness to pay-willingness to accept gap, the” endowment effect,” subject misconceptions, and experimental procedures for eliciting valuations. *The American Economic Review*, 95(3):530–545, 2005.
- John G Riley and William F Samuelson. Optimal auctions. *The American Economic Review*, 71(3):381–392, 1981.
- Ben Riley-Smith. Donald trump 2020: Us president launches re-election campaign in orlando vowing political ‘earthquake’. *The Telegraph*, 2019. URL <https://www.telegraph.co.uk/news/2019/06/19/donald-trump-2020-us-president-launches-re-election-campaign/>.
- Lee Ross. Negotiation and conflict resolution. *Barriers to Conflict Resolution*, page 26, 1995.

- Jean-Jacques Rousseau. A discourse on inequality, translated from french by maurice cranston, 1755.
- Thomas H Sawyer. Private golf clubs: Freedom of expression and the right to privacy. *Marq. Sports LJ*, 3:187, 1992.
- Jason F Shogren, Seung Y Shin, Dermot J Hayes, and James B Kliebenstein. Resolving differences in willingness to pay and willingness to accept. *The American Economic Review*, pages 255–270, 1994.
- Stephanie M Smith and Ian Krajbich. Gaze amplifies value in decision making. *Psychological Science*, 30(1):116–128, 2019.
- Mary Ellen Stratthaus. Flaw in the jewel: housing discrimination against jews in la jolla, california. *American Jewish History*, 84(3):189–219, 1996.
- Richard Thaler. Toward a positive theory of consumer choice. *Journal of Economic Behavior & Organization*, 1(1):39–60, 1980.
- R Blythe Towal, Milica Mormann, and Christof Koch. Simultaneous modeling of visual saliency and value computation improves predictions of economic choice. *Proceedings of the National Academy of Sciences*, 110(40):E3858–E3867, 2013.
- Hal R Varian. Two problems in the theory of fairness. *Journal of Public Economics*, 5 (3-4):249–260, 1976.
- Thorstein Veblen. The theory of the leisure class. reprint 1965, 1899.
- C Von Wilpert. Supreme streetwear’s insane success. *Sumo Blog*, 2019. URL https://sumo.com/stories/supreme-marketing?utm_source=growrevenue.io&utm_medium=Community&utm_campaign=growrevenue.
- Martin L Weitzman. Is the price system or rationing more effective in getting a commodity to those who need it most? *The Bell Journal of Economics*, pages 517–524, 1977.
- Robert B Wilson. Communications to the editor—competitive bidding with disparate information. *Management Science*, 15(7):446–452, 1969.
- Jenna Wortham. To pique interest, start-ups try a digital velvet rope. *New York Times*, 2011. URL <https://www.nytimes.com/2011/07/18/technology/start-ups-try-to-beckon-users-by-invitation-only.html?r=1>.
- Daniel Dylan Wray. Studio 54 – the story of the most famous club in the world, remembered by the man who just had to buy it.

Loud and Quiet, 1989. URL <https://www.loudandquiet.com/interview/studio-54-the-story-of-the-most-famous-club-in-the-world-remembered-by-the-man-who-just-had-to-buy>

8 Appendix: Proofs

Proof of Corollary 1. Suppose person i has the object. If $v_i < v_j$, trade takes place if and only if (iff) $(1 - \alpha)v_i + \alpha v_j + \varepsilon < v_j - \varepsilon$. The probability that this inequality is satisfied is strictly decreasing in α and becomes zero for any α sufficiently high, but still smaller than one. If $v_i > v_j$, trade never occurs for any $\alpha \leq 1$. ■

Proof of Corollary 2. Consider $p \in (0, \bar{v})$. If $\alpha = 0$, the seller accepts iff $p \geq v_s$ and the buyer accepts iff $p \leq v_b$. Consider $\alpha > 0$. If the seller says no, the buyer's payoff is independent of whether she says yes or no. Furthermore, the buyer's strategy must be monotone since, holding the seller's strategy constant, the buyer's utility from the object is increasing in v_b given any $\alpha < 1$. If the buyer says 'no,' this affects the seller's utility, but whether the seller says 'yes' or 'no,' has no bearing on the seller's payoff. Since the buyer's strategy is monotone, the seller's value of keeping the object must be increasing in v_s given any $\alpha < 1$. Let then $\underline{v}_b$ be the inf buyer type saying yes, and let $\bar{v}_s$ be the sup seller type saying yes. Suppose that $\underline{v}_b < \bar{v}_s$. Then $\bar{v}_s < p$ for any $\alpha > 0$. It follows that $\underline{v}_b < p$. This is, however impossible. Conditional on trade the buyer's expected utility would then be $\underline{v}_b + \alpha z - p$, where z is bounded by $\bar{v}_s - \underline{v}_b$. Hence $\underline{v}_b + \alpha z - p < (1 - \alpha)\underline{v}_b + \alpha\bar{v}_s - p < (1 - \alpha)(\underline{v}_b - \bar{v}_s) < 0$. Suppose that $\underline{v}_b = \bar{v}_s$. Then, the seller has an incentive to deviate as long as $\alpha > 0$. Hence, $\underline{v}_b > \bar{v}_s$ must hold. ■

Proof of Corollary 3. An increase in C decreases the utility boost from holding the cup and leaves the utility from holding a pen unaffected. The logic with respect to a decrease in P is analogous. ■

Proof of Proposition 1. Let the seller-optimal posted price be p^* and the associated expected profit, per buyer, be V . If $\alpha = 0$, the seller-optimal mechanism is charging p^* to all buyers with an expected overall profit of MV , e.g., [Harris and Raviv \(1981\)](#). Consider $\alpha > 0$. Excluding a single buyer leads to an expected loss of V . Holding p^* constant, the probability that any given remaining buyer buys is now raised by some $q > 0$, increasing in α , since $p^* < \bar{v}$ must hold. The gain in the expected profit is $(M - 1)qp_M$. In turn, there exists M_α^* such that if $M > M_\alpha^*$, then $(M - 1)qp_M > V$. By revealed preference, re-optimizing the uniform price can only further raise the seller's expected profit. The same argument holds if exclusion per buyer binds only probabilistically with a strictly positive probability for any given buyer type v . ■

Proof of Corollary 4. Suppose that $v^h > p^*$. The expected loss from exclusion is now $V' \leq V$, the expected gain is $(M - 1)q'p^*$ where $q' > 0$ still holds since there is a

strictly positive measure of types below p^* who are now willing to buy at p^* , but did not buy without exclusion, and the previous argument applies. Suppose that $p^* \geq v^h$. Then there is no loss from v^h -constrained random exclusion, since only types who were excluded by the price p^* are effectively targeted, but the WTP of types strictly below v^h increases. ■

Proof of Proposition 2. Let $b(v)$ be the symmetric monotone bidding strategy with $b(0) = 0$. Let's denote the payoff, maintaining equilibrium behavior by others, when type v pretends to be type z by $EU(v, z)$ and let $G(z)$ be the cdf of the highest intrinsic taste amongst the remaining $N - 1$ active bidders. In case of a downward deviation,

$$EU(v, z) = G(z)[v - b(z) + \alpha \int_v^{\bar{v}} (x - v)h(x)dx],$$

where $H(y)$ is the cdf of the highest intrinsic taste of the excluded bidders. In turn,

$$EU_z(v, z) = g(z)v - g(z)b(z) - G(z)b'(z) + g(z)\alpha \int_v^{\bar{v}} (x - v)h(x)dx$$

Implying the local optimality condition of:

$$b(v) = \frac{1}{G(v)} \int_0^v g(x)[x + \alpha K(x)]dx,$$

where, with a slight abuse of notation, $K(x) = \int_x^{\bar{v}} (y - x)h(y)dy$, with $K(x) \equiv 0$ if $K = 0$. To show that downward deviations are not profitable, consider $EU(v, z) - EU(v, v)$ given $z < v$. Note first that this difference can be written as:

$$(G(z) - G(v))(v + \alpha K(v)) + \int_z^v g(x)(x + \alpha K(x))dx.$$

Substituting terms and integrating by parts, the difference is:

$$(1 - \alpha)G(z)(v - z) - (1 - \alpha) \int_z^v G(x)dx + \alpha \left[\int_z^v G(z)H(x)dx - \int_z^v G(x)H(x)dx \right] < 0.$$

Consider now $z > v$. Here,

$$\begin{aligned} EU(v, z) &= G(z)(v - b(z)) + \\ &\quad \alpha \int_v^z (x - v)(g(x)H(x) + G(x)h(x))dx + \alpha G(z) \int_v^{\bar{v}} (x - v)h(x)dx. \end{aligned}$$

$EU_z(v, z)$, evaluated at $z = v$, we again get:

$$EU_z(v, v) = g(v)v - g(v)b(v) - G(v)b'(v) + \alpha g(v) \int_v^{\bar{v}} (x - v)h(x)dx.$$

Hence, the local necessary condition for optimality is the same. Consider now the difference $EU(v, z) - EU(v, v)$:

$$\begin{aligned} & G(z)(v - z) + \int_v^z G(x)dx - \alpha \int_v^z K(x)g(x)dx + \\ & + \alpha \left[\int_v^z (x - v)(h(x)G(x) + H(x)g(x))dx + \right. \\ & \left. + G(z) \int_v^{\bar{v}} (x - v)h(x)dx - G(v) \int_v^{\bar{v}} (x - v)h(x)dx \right]. \end{aligned}$$

Note that $\int_v^z K(x)g(x)dx = K(z)G(z) - K(v)G(v) + \int_v^z (1 - H(x))G(x)dx$, since $K'(x) = -(1 - H(x))$. Hence, the above can be further written as:

$$\begin{aligned} & (1 - \alpha)[G(z)(v - z) + \int_v^z G(x)dx] + \\ & + \alpha[-K(z)G(z) + K(v)G(v) + \int_v^z H(x)G(x)dx + \\ & + \int_v^z (x - v)(h(x)G(x) + H(x)g(x))dx + \\ & + G(z) \int_z^{\bar{v}} (x - v)h(x)dx - G(v) \int_v^{\bar{v}} (x - v)h(x)dx - G(z)(z - v)]. \end{aligned}$$

Simplifying terms, $K(v)G(v) = G(v) \int_v^{\bar{v}} (x - v)h(x)dx$, and $G(z) \int_z^{\bar{v}} (x - v)h(x)dx = G(z)K(z) + G(z)(z - v)(1 - H(z))$, the part inside the second square brackets can then be written as:

$$\begin{aligned} & \int_v^z H(x)G(x)dx + \int_v^z (x - v)(h(x)G(x) + H(x)g(x))dx + \\ & + G(z)(z - v)(1 - H(z)) - G(z)(z - v) \end{aligned}$$

which is non-positive, and the overall difference is negative.

It follows that if $\alpha = 0$, bidding is independent of K . Note also that for $K > 0$, $b(v)$ can be written as

$$v + \alpha K(v) - \int_0^v \frac{G(x)}{G(v)} [1 - \alpha(1 - H(x))]dx,$$

where $1 - \alpha(1 - H(x)) > 0$ which is independent of N . Hence, holding K constant, bids are increasing in N . Similarly, holding N , constant, bids are increasing in K since $K(v)$ increases in K and $H(x)$ is decreasing in K for any $x < 1$. ■

Proof of Proposition 3. From the proof of Proposition 2 it follows that:

$$b(v) = (1 - \alpha) \frac{N-1}{N} v + \alpha \frac{K}{K+1} + \alpha (1 - \frac{K}{K+1}) (\frac{N-1}{K+N}) v^{K+1},$$

and $\Pi(M, K)$ is thus:

$$(1 - \alpha) \frac{M-K-1}{M-K+1} + \alpha [\frac{K}{K+1} + (1 - \frac{K}{K+1}) (\frac{M-K-1}{M}) \frac{M-K}{M+1}].$$

Ignoring integer constraints, consider $\Pi_K(M, K)$. Note that $\Pi_K(M, K) < 0$ if $\alpha = 0$ and $\lim_{\alpha \rightarrow 1} \Pi_K(M, K) > 0$, with $\Pi_K(M, K)$ increasing in α and $\Pi_{K,K}(M, K) < 0$. It follows, that there exists α_M such that iff $\alpha > \alpha_M$, then $\Pi(M, K)$ is globally increasing in K . Simple calculations show that $\Pi(M, 1) > \Pi(M, 0)$ iff $\alpha > \alpha^*$. If $\alpha > \alpha^*$, there then exists $K_{M,\alpha}$ such that $\Pi(M, K) > \Pi(M, 0)$ if $K \leq K_{M,\alpha}$ where $K_{M,\alpha}$ is increasing in α . In turn, for $\alpha \in (\alpha^*, \alpha_M)$, $\Pi(M, K)$ is inverse U -shaped in K . Finally, to show that $K_{M,\alpha}$ is increasing in M , note that $\text{sign}\{\Pi(M, K) - \Pi(M, 0)\} = \text{sign}\{\alpha(1-K) + M(3\alpha-2)\}$. In turn, if for a given α and K , this difference is positive at some M , it is also positive for any $M' > M$. Corollary 5 follows from the above. ■

Proof of Proposition 4. Under lowest exclusion, the winner of the auction is the person with the highest consumption utility and her expectation of the second highest active valuation is unaffected. In turn, for any fixed M , the bidding strategy, hence, the expected revenue is the same as under no exclusion, $\Pi(M, 0) = \Pi_{Low}(M, K)$ for any α and any K . Fix M and K . Let $\alpha = 0$. Under random exclusion the bidding strategy is $b(v) = v - F(v)^{1-M+K} \int_0^v F(x)^{M-K-1} dx$, under lowest exclusion it is $b_{Low}(v) = v - F(v)^{1-M} \int_0^v F(x)^{M-1} dx$. ■

Proof of Proposition 5. If $\alpha = 0$, then $b_i(v_i) = v_i$, is the unique equilibrium since a person's payoff is not affected by the strategies of others. Note that for any given p , holding the other players' strategies constant, the payoff from 'winning' the object is strictly increasing in v_i as long as $\alpha < 1$. Hence, $b_i(v)$ can not be decreasing in v_i since, holding the strategies of the other players constant, this would violate incentive compatibility. Hence, b_i must be monotone increasing since the bid does not directly affect the price paid only the probability of winning. In turn, given symmetric strategies, the superiority boost can only come from the excluded ones and in equilibrium $b_i = v_i + \alpha E \max_{j \in K} \{v_j - v_i, 0\}$ holds. ■