RICCARDO CAMBONI

University of Padua

LUCA CORAZZINI

University of Milan Bicocca and Masaryk University

STEFANO GALAVOTTI

University of Bari and Masaryk University

PAOLA VALBONESI

University of Padua

PROCURING INNOVATION THROUGH BID PREFERENCE: AN EXPERIMENT

January 2026

Marco Fanno Working Papers – 331

Dipartimento di Scienze Economiche e Aziendali 'Marco Fanno'

Università degli Studi di Padova

Procuring Innovation Through Bid Preference: An Experiment*

Riccardo Camboni† Luca Corazzini‡ Stefano Galavotti§ Paola Valbonesi¶

Abstract

Procuring innovation involves motivating research efforts to design an innovative project and then implementing that project in a cost effective manner. In this paper, we study how to combine these two goals in a single procurement mechanism. To this end, we run an experiment in which subjects make costly investments in innovation, and then the most innovative project is selected and awarded through a first-price procurement auction. In one of our two treatments, the bidder who proposes the most innovative project in the first stage receives a bid preference in the auction.

The theory predicts that the bid preference scheme should indeed boost innovation, but at the cost of a higher awarding price in the auction. Our experimental results confirm the former prediction; however, we find that the bid preference does not result in an extra-payment for the buyer. To explain this puzzle, we search for the presence of a sunk cost fallacy in individuals' behavior. We provide robust evidence of this bias from two sources. First, in a subsequent Becker-DeGroot-Marschak task played by subjects after the main experimental task, we show that incurring a sunk entry cost causally increases subsequent offers. Second, a Quantal Response Equilibrium model of our procurement of innovation game confirms the presence of a significant *reverse* sunk cost fallacy. However, the introduction of the bid preference appears to weaken the strength of this bias, which ultimately explains our puzzling result regarding the awarding price.

JEL classification: C90; D44; H57.
Keywords: Procurement of Innovation; Procurement Auctions; Sunk Cost Fallacy; Quantal Response Equilibrium; Laboratory Experiment.

*We thank Elisabetta Iossa, Massimo Motta, Mirco Tonin, Patrick Rey, and participants in the seminar organized by the Free University of Bozen, the ASSET 2025 Conference in Rabat for valuable comments and suggestions. We are also indebted with Andrea Albarea and the whole staff of the VERA–LabEx for helping us run the experiment. This research was generously supported by the grant"Progetto di Eccellenza, Fondazione Cariparo, bando 2017 - VALB_ECCE 18_01", P.I.: Valbonesi. This study was also funded by the Europe Union - NextGenerationEU, in the framework of the GRINS - Growing Resilient, INclusive and Sustainable project (GRINS PE00000018 - CUP B13D22001400004). The experimental protocol received ethics approval from the Ethical Committee of the University of Venice "Ca' Foscari" (10/2022).

†**Corresponding Author.** Department of Economics and Management, University of Padua, Italy.

‡Department of Economics, Management and Statistics, University of Milan Bicocca, Italy; and Masaryk University Experimental Economics Laboratory, Czech Republic.

§Department of Economics, Management and Corporate Law, University of Bari, Italy; and Masaryk University Experimental Economics Laboratory, Czech Republic.

¶Department of Economics and Management, University of Padua, Italy.

1 Introduction

In OECD countries, public procurement accounts for around 13% of GDP (OECD, 2017a). The recent and widespread tightening of government budgets has increased interest in fostering innovation through public procurement as a way to improve the productivity of public spending and to promote competitiveness and growth. Successful examples of government involvement as both a *user-demander* and a *financer-producer* of innovation have further highlighted this potential.[1] It is thus not surprising that almost 80% of OECD countries currently adopt measures to support innovation through procurement (OECD, 2017).

The procurement of valuable innovation involves two fundamental steps. The first step is to motivate R&D efforts and investments to develop ideas and design innovative solutions. The second step is implementation (or commercialization): ideas and innovative solutions must then be translated into products and services efficiently and cost-effectively. Achieving both objectives simultaneously is a delicate task, as it requires addressing the well-known trade-off between incentivizing innovation and promoting competition.

One solution to this trade-off in a public procurement setting is to reward the firm that proposes the best innovation with an advantage at the time of awarding the contract for project execution.[2] In this paper, we study this practice through a laboratory experiment. Specifically, in the first stage of the experiment, which we refer to as the *innovation stage*, suppliers compete to submit innovative proposals that involve costly investments (or research efforts). The most innovative proposals are then selected and announced, and in the second stage – the *auction stage* – suppliers compete for the right to execute it in a first-price, private value procurement auction. The crucial element is that the supplier that submitted the most innovative proposal in the innovation stage can (treatment B) or cannot (treatment S) be granted an advantage (bid preference) over the other supplier in the auction stage.

Our study is the first attempt to experimentally investigate a setting in which the results of both stages of the innovation procurement process are endogenously determined. In particular, our objective is to assess whether and to what extent the introduction of a bid preference in the auction stage affects incentives to innovate in the first stage and the overall performance of this procurement of innovation process. A notable feature of our setup is that investments in innovation are made upfront and, thus, they are sunk when suppliers compete in the auction. Given the vast empirical evidence showing that both laboratory subjects and real-world agents are often prone to a *sunk cost fallacy*, it is natural to investigate whether such bias is present in our setting and, if so, what its implications are for bidding behavior and the overall procurement outcome.

Our main findings can be summarized as follows. First, the introduction of a bid preference in the auction stage does indeed stimulate higher investment in the innovation

[1] For example, the research commissioned by the US government in the 1960s to develop computer network communications, which laid the groundwork for the Internet, and the US supercomputing procurement program, which led to the development of the first processor on a chip.

[2] In the literature, this practice is referred to by various terms, including “bid preference”, “bid subsidy”, “bid credit”, “bonus system”, “bid discount”, and “price preference”.

stage, as predicted by theory, although the observed effect is weaker than expected. Second, when looking at the effects of the first stage on the second, we observe that suppliers, upon bidding in the auction, respond to the outcome of the innovation stage beyond what is predicted by the theory. Theoretically, there should be no other link between the outcome of the first stage and the bidding behavior, besides the fact that, when there is a bid preference, the supplier that submitted the most innovative proposal, having an advantage in the auction, should bid less aggressively: this softens competition, thus increasing the expected price paid by the buyer. Although we do indeed observe such bid spread in the treatment with bid preference, as expected, we do not see any significant difference in the expected price paid by the buyer with respect to a standard auction.

To explain this discrepancy, we show that subjects in both treatments exhibit a sunk cost fallacy: their bids in the auction stage increase with the investment cost incurred in the innovation stage. Interestingly, this bias operates in the reverse direction of what is typically observed. Rather than responding to a large sunk investment by reducing their bids in order to win and justify the initial expense, subjects do the opposite: they increase their bids, lowering their chances of winning in exchange for a higher potential profit. A Becker-DeGroot-Marschak task (with entry) played after our main experiment provides further and clean evidence that subjects are indeed prone to this bias.

Although this reverse sunk cost fallacy is observed in both treatments, our structural analysis shows that it is stronger when the bid preference is absent, and this helps explain why there is no significant difference between treatments in the price paid by the buyer to the winning bidder. We conjecture that this difference in the strength of the bias is related to the cognitive load associated with each treatment. In the setting without bid preference, the contract is awarded through a standard first-price auction, and subjects appear to focus more on the sunk investment they have just made, bidding higher to secure a satisfactory payoff in case of winning. In contrast, the setting involving the bid preference is cognitively more demanding, as subjects must take into account that their bids are treated differently in the determination of the winner. This complex awarding rule is likely to capture most of their attention in the bidding stage, making the sunk cost less salient. In other words, the higher cognitive load induced by the introduction of the bid preference may partially de-bias participants, reducing the impact of the sunk cost fallacy. As a result, the expected drawback of the bid preference scheme -— namely, a higher price paid by the buyer — does not materialize.

The remainder of the paper is structured as follows. Section 2 reviews the relevant literature. Section 3 describes the experimental design. Section 4 presents the theoretical predictions. Section 5 reports and discusses the results of the experiment, including a structural analysis. Section 6 concludes.

2 Related literature

It is well known that innovation is a crucial engine for productivity and economic growth (see, e.g., Griffith et al., 2004). However, market economies tend to underprovide R&D (see, e.g., Arrow, 1972) due to the public good nature of innovation.

These facts have supported the development of demand-side innovation policies and, in particular, have drawn attention to the potential of governments' substantial purchasing power to promote innovation (Cabral et al., 2006), giving rise to a growing body of literature on the "public procurement of innovation" (see the surveys by Cabral et al., 2006, Iossa et al., 2018, and Chiappinelli et al., 2025). On the one hand, studies have estimated the impact of public procurement on private R&D investment (Slavtchev and Wiederhold, 2016), on the long-term profitability of businesses (Howell, 2017), and on the determinants of successful R&D contests (Bhattacharya, 2021, and Decarolis et al., 2021). On the other hand, in a number of recent reports and official initiatives by the European Commission (e.g., the "Guidance on Innovation Procurement" delivered in 2021) and the OECD (OECD, 2017), the potential of public procurement of innovation has been emphasized and encouraged in its application.

This debate has naturally drawn attention to the design of appropriate incentives to boost innovation through public procurement. Several theoretical contributions (Bennett and Iossa, 2006; Chen and Chiu, 2010; Hoppe and Schmitz, 2013; Iossa and Martimort, 2012) and one experiment (Hoppe et al., 2013) have emphasized the innovation potential associated with public-private partnerships, intended as the bundling of project design and implementation tasks under a single contractor or consortium. These papers focus mainly on post-contract (and non-verifiable) cost-reducing innovations, and notably do not address the problem of optimal selection among multiple contractors. In such settings, bundling encourages innovation by allowing the consortium to internalize its benefits.

When, instead, several potential suppliers are involved and the challenge is to incentivize research efforts prior to contracting, assigning both tasks under a single contractor is not necessarily optimal. A contractor selected on the basis of a project's innovativeness might not be the most efficient one to implement it. This more general setting is considered by Che et al. (2021). In their framework, firms develop and submit research ideas whose value is observable and positively influenced by their research efforts. In the subsequent production stage, each firm submits a bid for each proposed research idea. The contracting authority then selects both the project to be implemented and the firm to carry it out, thus implying that the implementing firm may differ from the proposing firm. They show that the optimal mechanism involves a trade-off between incentivizing R&D and achieving allocative efficiency at the implementation stage. Specifically, to encourage research effort, the awarding rule is distorted in favor of firms that submitted more innovative proposals, meaning that project execution may not be awarded to the lowest-cost bidder. The optimal mechanism thus illustrates the effectiveness of production rights, such as bid preferences, as a tool to incentivize innovation, even if this may increase implementation costs.[3]

Our paper is closely related to Che et al. (2021), as we implement in the lab a stylized version of their optimal mechanism to explore how this trade-off is resolved by experimental subjects. In particular, compared to their model, we consider a simplified setting with only two suppliers, deterministic innovation, project selection occurring before the auction, and a first-price auction in which the innovating firm may receive a (fixed) preferential treatment instead of the optimal mechanism.

[3]See Hodges and Dellacha (2007) for a discussion of this type of mechanism and the various forms used in practice. Recently, bid preference schemes have become common in Green Public Procurement, favoring firms that adopt green technologies (see Chiappinelli and Seres, 2024).

Our paper also indirectly contributes to the literature on auctions in which bids or bidders are treated asymmetrically, for various reasons. McAfee and McMillan (1989), in the context of international trade, show that if foreign firms have cost advantages, governments should discriminate in favor of domestic firms to stimulate competition and minimize expected procurement costs.[4] Marion (2007) empirically studies bid preferences for small businesses in California highway procurement auctions and finds that procurement costs were 3.8% higher when bid preferences were applied. He attributes this increase to decreased participation by large firms in subsidy auctions. Krasnokutskaya and Seim (2011) estimate a structural bidding model using a dataset on Californian highway subsidy procurement auctions, and employ simulations to evaluate alternative preference policies. Athey et al. (2013) investigate set-asides in US Forest Service timber auctions and use the results to estimate the impact of bid preferences on the same dataset. Corns and Schotter (1999) contrast these empirical findings with experimental evidence, showing that bid preference rules can increase both minority representation and cost-effectiveness if the preference level is appropriately calibrated. However, in order to make such a choice, the cost distributions of the bidders must be known.

Finally, our paper contributes to the literature on the sunk cost fallacy, a well-known psychological phenomenon whereby individuals violate the principle that past sunk costs should be irrelevant to current decisions. The conventional view holds that the sunk cost fallacy manifests itself as a "tendency to continue an endeavor once an investment in money, effort, or time has been made" (Arkes and Blumer, 1985, p. 124), even when continuing is not optimal. This view is supported by a number of hypothetical surveys (see, e.g., Arkes and Blumer, 1985, Arkes and Ayton, 1999, Strough et al., 2008, Molden and Hui, 2011) and some anecdotal evidence.[5] Various behavioral explanations have been proposed, including waste aversion (Arkes and Blumer, 1985; Arkes and Ayton, 1999), self-justification (Staw, 1976), prospect theory (Whyte, 1986), mental accounting (Ho et al., 2018), limited memory (Baliga and Ely, 2011), and self-control commitment devices (Eswaran and Neary, 2016; Hartig, 2017; Hong et al., 2019). Despite this conventional view, there is little conclusive evidence from incentivized experiments that individuals actually fall prey to this bias. Friedman et al. (2007) find a surprising small effect, and in Phillips et al. (1991), only one-quarter of participants behave in accordance with the bias. More importantly, several experiments have documented a *reverse* sunk cost effect: the greater the resources already invested, the less committed subjects become to continuing the project. For example, after a large initial investment, subjects often quit the project earlier than is optimal. This reverse sunk cost fallacy is clearly documented by Negrini et al. (2022) and is also observed in a few participants in Phillips et al. (1991) and in the benchmark treatment of Baliga and Ely (2011).[6] In this respect, by finding a positive rela-

[4] They suggest a rough rule of thumb to determine the appropriate preference level: use one-third of the percentage difference between the means of the cost distributions of the two types of firms as the price preference. This rule assumes uniform cost distributions.

[5] For instance, the sunk cost fallacy is often invoked to explain why politicians continue over-budget public projects or why firms keep investing in unprofitable ventures (hence the term *Concorde effect*). Augenblick (2016) explains bid escalation in penny auctions as an instance of sunk cost fallacy.

[6] In Baliga and Ely (2011), the reverse sunk cost fallacy is referred to as the *pro-rata effect*. Interestingly, they observe this bias in the benchmark treatment, where, according to their behavioral model, no bias

tionship between investment costs in the Innovation stage and bids in the Auction stage, our paper provides further evidence that the sunk cost fallacy may also manifest itself in this reverse direction.

3 Experimental design

The experiment consists of two consecutive games, described in the following. The first game, called *"Procurement of Innovation Game"*, is meant to address our main research question as it embodies the two fundamental steps of innovation – R&D and implementation – in a single procurement setting. In the Procurement of Innovation Game, the buyer solicits the submission of proposals, selects the most innovative one, and – provided it is sufficiently satisfactory – proceeds to its execution by auctioning it off among the same firms that submitted the proposals in the first place. The experimental subjects play the role of potential suppliers who strategically interact to yield innovative proposals and then eventually sell the most innovative project to the buyer.

One crucial element of the Procurement of Innovation Game is that the investment in innovation is made upfront and thus represents a sunk cost for the suppliers. Informed by the literature on the sunk cost fallacy, we suspected that this bias may alter suppliers' behavior and eventually affect the outcome of the procurement process.

The second game of our experiment, called *"Market Game"*, is specifically designed to better isolate the sunk cost fallacy. Like the Procurement of Innovation Game, the Market game involves two stages and a sunk investment in the first; however, it is immune from strategic interactions, as each subject simply plays the role of a seller who attempts to sell a hypothetical object to an automated buyer.

The details of the Procurement of Innovation Game and of the Market Game are described in the following.

3.1 Procurement of Innovation Game

The Procurement of Innovation Game (*PING*, henceforth) consists of two sequential stages, the *innovation stage* and the *auction stage*, and involves two active players, numbered 1 and 2, who act as suppliers that invest in innovation and then compete for the right to execute the most innovative project for a fictitious automated buyer. The idea is that the innovation stage determines the characteristics – specifically, the degree of innovativeness – of the project that is later auctioned off in the auction stage.

In the innovation stage, the two suppliers simultaneously choose a level of innovation, which must be an integer in the set $\{0, 1, 2, 3, 4\}$. Let i_1 and i_2 denote the innovation levels chosen by suppliers 1 and 2, respectively. Innovation is costly: choosing an innovation level i entails an investment cost $k(i) = i^2$. The supplier who chooses the higher innovation level is designated the *innovator* (denoted by I), while the other becomes the *challenger* (denoted by C); we let $i_I = \max\{i_1, i_2\}$ denote the innovation level chosen by the innovator.[7] Crucially, i_I defines the degree of innovativeness of the project that is then auctioned

should be present.

[7] Ties in innovation levels are broken randomly.

off.

At the end of the innovation stage and before entering the auction stage, suppliers learn their role, I or C, and the value of i_I. In the auction stage, the two suppliers compete for the right to execute the project: they privately observe their own production costs, c_I and c_C, that are independently drawn from a discrete uniform distribution over the integers from 0 to 100, and simultaneously submit bids, denoted b_I and b_C.[8] Bids must be integers and cannot exceed the reserve price, set at 135.

The winner of the auction is determined by one of the two allocation rules described below. However, there is a commonly known probability that the auction will be revoked ex-post. Specifically, the probability that the auction is *not* revoked is:

$$q(i_I) = \frac{1 + i_I}{5},$$

which increases with the innovation level of the project, as determined in the innovation stage.[9] If the auction is revoked, the project is not executed. If the auction is not revoked, the project is executed: the winner is paid her bid and incurs her production cost. Note that the two stages of the game are linked in two respects: (i) the innovation stage determines the roles, either I or C, that the two suppliers have in the auction stage; (ii) the innovation level i_I affects the probability that the project is implemented.

The suppliers' monetary payoffs are given by the (algebraic) sum of their initial endowment of 16 and the gains/losses obtained in the two stages. In the innovation stage, each supplier bears the cost associated with her innovation choice. In the auction stage, the winner's payoff is the difference between her bid and her production cost, if the auction is not revoked, it is zero otherwise; the loser's payoff is always zero. In particular, the investment costs are sunk: they are incurred regardless of whether the project is executed or not.

We implemented two treatments that differ in how the winner of the auction is determined. In treatment *S* (Standard), the lowest bid wins. In treatment *B* (bid preference), I is given an advantage over C: in particular, for the sake of the winner's determination, the bid of C is increased by a multiplicative factor equal to 1.5. In other words, C wins the auction only if $1.5 \times b_C < b_I$. Notice that the multiplicative factor 1.5 does not affect the payment to the winner of the auction. In other words, if $1.5 \times b_C < b_I$, C wins the auction and is then paid her genuine bid b_C, not her augmented bid ($1.5 \times b_C$).[10]

3.2 Market Game

The Market Game (*MG* henceforth) includes two sequential stages like *PING*, but lacks strategic interaction between subjects. Specifically, each subject receives an initial endowment of 9. In the first stage, the subject is randomly assigned an entry cost $e \in$

[8] The cost distribution is the same for both roles (i.e., the innovator has no cost advantage). This captures the pure public good nature of the innovation.

[9] In the experiment, revoking the auction is determined *after* the auction stage to gather more data. Conceptually, however, it corresponds to a decision made by the buyer upon receiving and evaluating proposals. Importantly, the execution decision is stochastic but independent of the auction stage.

[10] As in the innovation stage, ties in the auction are broken randomly.

$\{0,1,2,3,4,5,6,7,8,9\}$, and then chooses whether to *Enter* or *Not*. If the subject does not enter, the game ends and the subject's payoff is equal to the initial endowment of 9. If, instead, the subject decides to enter, she bears the entry cost and proceeds to the next stage.

In the second stage, the subject is assigned a production cost, randomly drawn from the set of integers between 0 and 15. Given the production cost, the subject chooses a selling offer for a hypothetical object that must be an integer number between 0 and 30. This is then compared to a randomly generated buying offer, also an integer between 0 and 30. If the buying offer is greater than or equal to the selling offer, the object is sold: the subject is paid the buying offer and incurs the production cost. Otherwise, no trade occurs, and the subject earns nothing in this stage.[11]

3.3 Procedures

Upon arrival, subjects were randomly assigned to computer terminals. At the beginning of the session, subjects were told that they would have participated in two consecutive games (referred to as "experiments" in the instructions) and that their final earnings would have been determined according to the results of one game only, randomly selected at the end of the experiment. The instructions for each task were distributed and read aloud at the beginning of the corresponding game (instructions for treatment *B* of *PING* are included in the Online Appendix). Before each game, subjects answered control questions to ensure comprehension. When needed, questions were checked and explained privately. Each session included 15 repetitions of the first game and 10 of the second.[12]

At the beginning of *PING*, the computer randomly formed 4 rematching groups, each containing 6 subjects. The composition of the rematching groups remained unchanged throughout the game. In each period of *PING*, subjects were anonymously and randomly paired within these groups, with the rule that no subject faced the same opponent in two consecutive periods.[13] No rematching groups were needed in *MG* of the second phase, as subjects did not strategically interact with each other.

In *PING*, at the end of the innovation stage, each subject was shown her own choice (and cost), her assigned role (either I or C), the innovation level chosen by the innovator (i_I), the resulting probability of revocation, and her production cost. At the end of the auction stage, a summary of the results of both stages was displayed, including whether the auction was revoked and the subject's total earnings.

In *MG*, a subject who chose to enter and bear the corresponding cost was shown her decision (and entry cost), and her production cost; at the end of the game, the screen displayed the buying offer, the trade outcome, and the subject's earnings.

Before confirming their choices, subjects could simulate the consequences of their decisions on the screen: investment costs, auction outcomes, probability of revocation, trade possibility, and earnings.

[11] Note that the second stage of *MG* is a Becker-DeGroot-Marschak task (Becker et al., 1964).

[12] We use the terms repetition and period interchangeably.

[13] With 6-person groups, subjects interacted with the same opponent every 5 periods on average. While not a perfect stranger protocol, this left little room for repeated-game strategies. The rematching protocol was aimed at increasing independent observations and facilitating non-parametric robustness checks.

We ran 3 sessions for each of the 2 treatments of *PING*, each involving 24 subjects, thus producing 12 independent observations at the rematching group level. The second part (*MG*) was identical across sessions.

The experiment was conducted at VERALabEx, University of Venice, between January and March 2023. The participants were mainly undergraduate students, recruited through the VERA-Lab online system. The experiment was computerized using the *z-Tree* software (Fischbacher, 2007). Prices, costs, and earnings were denominated in points. At the end of each session, one of the two experimental games was randomly selected to determine the payments to the subjects, and the experimental points were converted into euros and paid privately in cash. In order to ensure the same expected payment for the two games, we used two different exchange rates: 33 points for 1 euro in *PING* and 12 points for 1 euro in *MG*.

On average, subjects earned approximately 15 euros (including a show-up fee of 4 euros) for sessions lasting 75 minutes, including instructions and payment. Before leaving, participants completed a short questionnaire on demographics and perceptions of the experimental task.

4 Theory and predictions

4.1 Procurement of Innovation Game (*PING*)

For the two treatments of *PING*, we derived testable predictions from the Perfect Bayesian Equilibria (PBE) of these games, assuming that both suppliers are risk-neutral, i.e. their payoffs correspond to their monetary earnings.[14] In light of the sequential structure of the games, we proceed backward, first analyzing the auction stage and then the innovation stage.

4.1.1 Auction stage

With risk-neutral and rational suppliers, bidding behavior in the auction stage is unaffected by the innovation level chosen in the innovation stage for two reasons: *(i)* upon bidding in the auction, the investment decisions made by the suppliers in the innovation stage are sunk; *(ii)* i_I, the innovation level chosen by the innovator affects $q(i_I)$ -— the probability that the auction outcome is implemented -— but a rational bidder should focus solely on her expected payoff from the auction, conditional on it being implemented. In treatment S, subjects' roles are irrelevant in the auction stage, and thereby, in equilibrium, the two suppliers behave symmetrically. In contrast, in treatment B, the roles are crucial, as the bid preference gives a competitive advantage to the innovator. As a result, the equilibrium is asymmetric.

Figure 1 shows the equilibrium bids in the auction stage for the two treatments, S and B.

[14] Equilibria are derived in the Appendix.

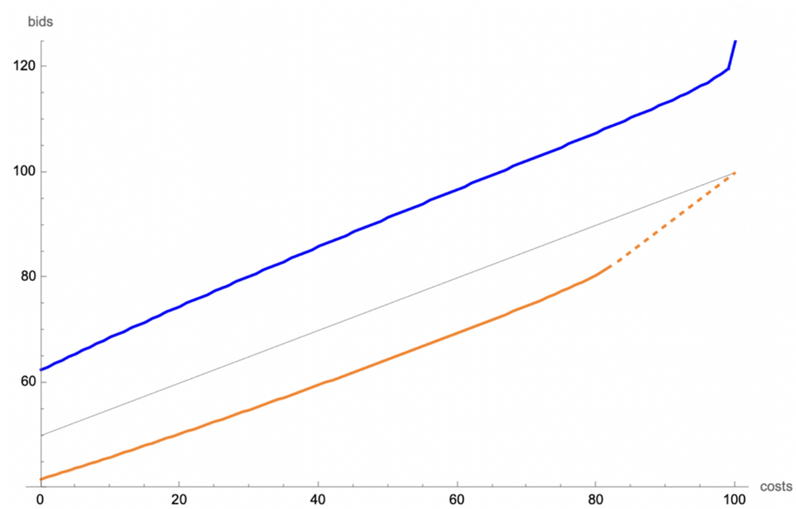

Figure 1: Equilibrium bidding functions in the auction stage for treatment S (light grey line) and B (blue line for the innovator, orange line for the challenger).

As is intuitive, in treatment B the bidding function of the challenger lies below that of the innovator. Moreover, certain types of challenger —- those with production costs greater than 84 -— have no chance of winning the auction, as their augmented bids exceed the innovator's maximum bid (125).[15] The (symmetric) equilibrium bidding function in treatment S lies between the innovator's and the challenger's bids in treatment B.

Table 1 summarizes the equilibrium results for the auction stage in terms of the expected winning bid, the expected production cost of the winner, the probability that the innovator wins and the expected payoffs (conditional on the auction not being revoked).

Table 1. Auction stage: Equilibrium analysis.

	$E[b_W]$	$E[c_W]$	$\Pr[I \text{ wins}]$	π_I	π_C
S	66.58	33.17	0.50	16.71	16.71
B	71.86	33.72	0.58	26.89	11.24

It is important to note that, in treatment B, despite the innovator bidding higher than the challenger, the former wins more often than the latter, thanks to the bid preference. As a result, the winning bid in treatment B tends to be higher than in treatment S. Moreover, since the innovator's and the challenger's production costs are drawn from the same distribution, the fact that, in treatment B, the innovator wins more frequently than the challenger implies that the project is sometimes inefficiently allocated —- i.e. the innovator sometimes wins even if her production cost is higher than that of the challenger. In contrast, in treatment S, the project is always allocated efficiently. As a result, the expected

[15] Note that in both treatments the highest equilibrium bid is well below the experimental reserve price of 135.

production cost of the winner in treatment B is larger than in treatment S.

4.1.2 Innovation stage

Let i_1 and i_2 denote the innovation levels chosen by the two suppliers, and let π_I and π_C be the equilibrium payoffs of the innovator and the challenger in the auction stage. If $i_1 > i_2$, then supplier 1 is the innovator, $i_I = \max\{i_1, i_2\} = i_1$, and supplier 1's total payoff is:[16]

$$\Pi_1 = 16 - k(i_1) + q(i_1)\pi_I.$$

If $i_1 = i_2$, the role of innovator is determined randomly, $i_I = \max\{i_1, i_2\} = i_1 = i_2$, and supplier 1's total payoff is

$$\Pi_1 = 16 - k(i_1) + q(i_1)\frac{1}{2}(\pi_I + \pi_C).$$

Finally, if $i_1 < i_2$, then supplier 1 is the challenger, $i_I = \max\{i_1, i_2\} = i_2$, and supplier 1's total payoff is

$$\Pi_1 = 16 - k(i_1) + q(i_2)\pi_C.$$

Note that the levels of π_I and π_C are equal in treatment S, but different in treatment B.

It turns out that the two treatments admit a symmetric equilibrium in mixed strategies:

- in treatment S, suppliers choose their innovation levels by randomizing over the set $\{0, 1, 2\}$, with (rounded) probabilities $\{0.30, 0.60, 0.10\}$;
- in treatment B, suppliers choose their innovation levels by randomizing over the entire set of possible investment levels $\{0, 1, 2, 3, 4\}$, with (rounded) probabilities $\{0.09, 0.13, 0.31, 0.12, 0.36\}$.

For treatment B, the above is the unique equilibrium of the game. On the other hand, treatment S also admits an asymmetric equilibrium in pure strategies where one supplier chooses $i = 2$ and the other chooses $i = 0$. However, our analysis will focus on the symmetric mixed equilibrium, which does not require coordination by bidders.

Table 2 presents the implications of the mixed equilibria described above in terms of innovation level, corresponding cost, and probability that the auction is executed.

Table 2. Innovation stage: Equilibrium analysis.

	$E[i]$	$E[i_I]$	$E[i_C]$	$E[k(i)]$	$E[k(i_I)]$	$E[k(i_C)]$	$E[q(i_I)]$
S	0.80	1.10	0.50	1.01	1.49	0.52	0.42
B	2.54	3.26	1.81	8.17	11.61	4.72	0.85

In treatment B, the average level of innovation chosen by the suppliers is three times higher than in treatment S. The same is true for the expected value of i_I, which corresponds to the level of innovation eventually incorporated into the project. As a result, the

[16] Since the game is symmetric, the payoff of supplier 2 is symmetric to that of supplier 1.

probability that the auction is executed in *B* is two times higher than in *S*. Observe, finally, that the average cost borne by the challenger is much higher in *B*: this has relevant welfare implications, as this cost represents a net social loss.

4.1.3 Testable predictions

The previous theoretical analysis is based on the standard hypothesis of fully rational, risk neutral players (and equilibrium behavior). This benchmark model yields clean predictions that can be tested through our experiment. First, from the point of view of the buyer, the bid preference scheme produces a trade-off: it triggers more innovation (which is beneficial for the buyer), but makes the buyer pay a higher price for the execution of the project.[17] Second, regardless of whether the bid preference is present or not, the behavior of bidders in the auction stage should not be affected by the level of investment made in the innovation stage, as this is totally sunk for the bidders. We summarize these results in the following predictions that represent our main null hypotheses to be tested in the experiment.

(P1) *The expected level of innovation is higher in treatment B than in treatment S.*

(P2) *The expected winning bid is higher in treatment B than in treatment S.*

(P3) *In the equilibrium of the auction stage, bidding functions are unaffected by the investment made in the innovation stage.*[18]

Although we are mainly interested in the effects of the bid preference from the buyer's viewpoint, the theoretical analysis highlights that a similar trade-off also affects the social welfare: the introduction of a bid preference, beyond generating more innovation (prediction P1), exacerbates the duplication of investment costs in the innovation stage and does not always yield an efficient allocation of the project to the lower cost supplier in the auction stage. We then add the following testable prediction regarding the inefficiency triggered by the bid preference scheme.

(P4) *The expected investment in innovation made by the challenger is greater in treatment B than in treatment S. Moreover, the expected production cost of the winner of the auction is higher in treatment B than in treatment S.*

[17] To determine the optimal solution to this trade-off, one should set out the payoff function of the buyer. We refrain from doing so, as any choice would be somewhat arbitrary. Our goal is just to examine if and to what extent this trade-off is confirmed by experimental data.

[18] Note that (P1), (P2) and (P3) correspond to the hypothesis (HP1), (HP4-1) and (HP3-1) pre-registered on AsPredicted.org (https://aspredicted.org/dg95u.pdf). In the pre-registration, we also included alternative behavioral hypotheses, based on the sunk cost fallacy, a bias that we suspected to possibly emerge in *PING*. These alternative hypotheses are discussed later in the paper (see Section 5.3). Notice also that the aforementioned pre-registration form also envisages two treatments that are the counterparts of treatments *S* and *B*, but where the innovation stage is suppressed and the roles (either *I* or *C*) are assigned randomly. In the current paper, we concentrate on treatments *S* and *B* only, as our focus here is on the use of a bid preference as an incentive to boost innovation. The results of the other two treatments will be analyzed in a separate paper centered on the study of the effect of the sunk cost fallacy in strategic contexts like auctions.

Finally, the analysis of the auction stage yields the following additional predictions concerning the shape of the equilibrium bidding functions.

(P5) *In the equilibrium of the auction stage:*

- *(a) bidding functions are strictly increasing in production costs;*
- *(b) in treatment B, for all production costs, the bid of the challenger is below that of the innovator, but the innovator wins more often than the challenger;*
- *(c) in treatment S, for all production costs, the bid of the challenger and the bid of the innovator are equal;*
- *(d) for all production costs, the bid of both bidders in treatment S is above the bid of the challenger in treatment B, and below the bid of the innovator in treatment B.*

4.2 Market Game

To determine the optimal decision of a risk neutral seller in *MG*, suppose that the seller has decided to enter, and has been assigned a production cost equal to γ. It is immediate to verify that it is always optimal for a subject to make a selling offer that is exactly equal to γ. Notice that this is the optimal selling offer regardless of the risk propensity of the seller.[19]

Going backward to the entry decision, since the expected payoff of a seller that enters is 8.9, it is optimal to enter if the entry cost is less than or equal to 8, while it is optimal not to enter only when the entry cost is equal to 9.

5 Experimental results

The experimental results are presented in three steps. First, we focus on *PING* to shed light on the above mentioned trade-offs affecting both the buyer's and the social welfare. Second, we look at the observed bidding behavior in *PING*, and we study the main determinants of the innovator's and the challenger's bids and how they respond to the introduction of the bid preference. Having documented an important discrepancy between the predicted and observed behavior, we finally investigate whether the sunk cost fallacy hypothesis is able to explain such a puzzle. In particular, after presenting prima-facie evidence in favor of the sunk cost fallacy hypothesis in the context of *MG*, we structurally isolate this behavioral component in *PING* by specifying and estimating a Quantal Response Equilibrium model that includes a sunk cost fallacy parameter.

[19]The argument is the same as the one showing that, in a second-price auction, it is a weakly dominant strategy to bid one's cost. In fact, like in a second-price auction, in *MG* the payment received by the seller in case of trade does not depend on the selling offer. Therefore, the selling offer only determines whether or not trade takes place. Now, for the seller it would be optimal to trade whenever the buying offer (which corresponds to the selling price if trade takes place) is greater than γ, while it would be optimal not to trade in the opposite case. A selling offer equal to γ ensures that the optimal outcome is implemented for all possible buying offers.

5.1 The trade-offs in buyer's and social welfare

Table 3 reports the descriptive statistics, distinguished by treatment, of the key outcome variables of *PING*: the investment levels in the innovation stage, the winning bid and the production cost of the winner in the auction stage.

We begin our analysis by focusing on the two determinants of the trade-off affecting the buyer's surplus: the presence of a bid preference stimulates innovation (prediction (P1)), but at the cost of a higher expected payment to the winning bidder (prediction (P2)).

Concerning innovation, results are in line with prediction (P1): the presence of a bid preference in treatment B generates a significantly higher innovation than in treatment S ($p = 0.022$).[20] However, while in treatment B we find no difference between the observed level of innovation and what theoretically predicted ($p = 0.105$), subjects in treatment S invest more than expected ($p < 0.001$). This over-investment partially, but not completely, closes the predicted gap between the two treatments. Finally, as a robustness check, we compare the innovation levels in B and S at the individual level, pooling data from all bids. The results confirm that innovation levels are higher in treatment B than in treatment S (see Appendix B, Table C2).

Concerning the winning bid, the key finding is that, contrary to what was predicted by (P2), the winning bids are not statistically different between treatment B and treatment S ($p = 0.933$). Furthermore, in both treatments, the winning bids are substantially lower than their theoretical values.[21]

Hence, the predicted trade-off affecting the buyer's surplus is not observed in the lab: while the bid preference effectively stimulates innovation (though the stimulus seems weaker than expected), it does not come at the cost of a higher price paid by the buyer to the winning bidder.

We now turn our attention to the trade-off that affects social welfare. As stated in prediction (P4), the positive impact of the bid preference on the realized level of innovation is accompanied by allocative inefficiencies: a higher duplication of investments in the innovation stage, and a higher production cost of the winning bidder.

Prediction (P4) is (weakly) confirmed in the data: the observed production cost of the winning supplier is somewhat higher in treatment B than in treatment S ($p = 0.072$). Notice also that this is larger than predicted in both treatments, indicating that subjects do not consistently behave as suggested by theory. In addition, the investment in innovation made by the challenger is significantly higher in treatment B than in treatment S ($p = 0.018$). Notice also that the innovation level chosen by the challenger is in line with its predicted level in treatment B, but significantly higher in treatment S, again reducing the expected gap in innovation between the two treatments.

[20] For the rest of the subsection, between-treatment differences in outcome variables are tested by using two-way linear random-effects models that account for both potential individual dependency over periods and dependency within rematching group. Models' outcomes are reported in Table C1 in Appendix B.

[21] The predicted winning bid is defined as follows: For treatment S, it is the theoretical bid of the supplier with the lowest production cost; for treatment B, where the lowest-cost supplier is not guaranteed to win, the prediction methodology accounts for the uncertainty at the beginning of the game about the innovator status. We calculated the theoretical winning bid under two equally likely scenarios — one for each bidder becoming the innovator — and defined the final predicted winning bid as the average of these two outcomes.

Table 3. Investment levels and bids in $PING$: descriptive statistics, by treatment.

Treatment	B	S
i_I	2.998	2.502
	(1.157)	(1.112)
$\widetilde{i_I}$	3.262	1.105
	(0)	(0)
$i_I - \widetilde{i_I}$	-0.271	1.386***
	(0.167)	(0.137)
i_C	1.561	0.983
	(1.365)	(1.118)
$\widetilde{i_C}$	1.831	0.500
	(0)	(0)
$i_C - \widetilde{i_C}$	-0.282	0.481***
	(0.212)	(0.112)
b^W	52.52	51.57
	(27.65)	(22.06)
$\widetilde{b^W}$	72.05	66.04
	(16.73)	(11.76)
$b^W - \widetilde{b^W}$	-19.391***	-14.111***
	(1.792)	(1.594)
c^W	38.89	35.46
	(26.81)	(25.47)
$\widetilde{c^W}$	34.31	32.09
	(25.02)	(23.53)
$c^W - \widetilde{c^W}$	4.574***	3.366***
	(0.661)	(0.462)
Observations	540	540

Notes. This table reports descriptive statistics (standard deviations in parentheses) of the auction stage in *PING*, distinguished by treatment. i_I and i_C are the *observed* average innovation levels of the innovator and the challenger, respectively. b^W and c^W are the *observed* average winning bid and production cost of the winning bidder, respectively. $\widetilde{i_I}$, $\widetilde{i_C}$, $\widetilde{b^W}$ and $\widetilde{c^W}$ are the corresponding *predicted* average values. The table also reports parametric estimates (standard errors in parentheses) from two-way linear random-effects models accounting for both potential individual dependency over periods and dependency within rematching group. The dependent variable is the difference between observed and predicted values, and the regression includes only the constant term as a regressor. Significance levels are denoted as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

We summarize the previous results in the following statements.

Result 1. Bid preference and buyer's surplus in *PING*. *The presence of the bid preference stimulates innovation, but it does not alter the winning bid in the auction stage.*

Result 2. Bid preference and social welfare in *PING*. *The presence of the bid preference leads to a higher investment cost borne by the challenger and a higher production cost of the winning supplier.*

5.2 Bidding behavior

Based on the previous analysis, we are led to reject prediction (P2): there is no significant difference in the winning bids in the auction stage between the two treatments. In addition, in both treatments, the winning bid seems to be lower than what was theoretically predicted. To shed light on these discrepancies, we examine in depth the bidding behavior in the auction stage of *PING*. Specifically, we are interested in studying the determinants of the innovator's and the challenger's bids, how they respond to the introduction of the bid preference, and whether they align with the predicted levels.

The descriptive statistics on the innovator's and the challenger's bids are presented in Table 4.

Table 4. Bids in the auction stage of *PING*: descriptive statistics, by treatment.

Treatment		*B*			*S*	
	All	*I*	*C*	All	*I*	*C*
b	67.80	75.70	59.90	66.75	67.43	66.07
	(31.42)	(30.18)	(30.67)	(26.81)	(27.15)	(26.48)
$\widetilde{b}$	78.88	90.78	66.99	74.35	74.58	74.12
	(20.12)	(15.89)	(16.56)	(14.43)	(14.36)	(14.50)
$(b-\widetilde{b})/\widetilde{b}$	-0.151	-0.175	-0.126	-0.122	-0.115	-0.129
	(0.345)	(0.297)	(0.387)	(0.285)	(0.292)	(0.278)
Observations	2160	1080	1080	2160	1080	1080

Notes. This table reports descriptive statistics (standard deviations in parentheses) of observed (b) and predicted bids ($\widetilde{b}$) in the auction stage of *PING*, as well as the corresponding levels of relative underbidding, defined as $(b-\widetilde{b})/\widetilde{b}$. Columns *I* and *C* refer to the bids of the innovator and the challenger, respectively.

A consistent pattern emerges: observed bids are systematically lower than their theoretical levels in both treatments and for both statuses. To better quantify this gap, we define a measure of *relative underbidding* as the difference between the observed and predicted bids, divided by the predicted level.[22] We find that bids are on average 15.1% lower

[22] Underbidding in our procurement setting where bidders are sellers corresponds to overbidding in standard auctions, namely a situation in which bidders (buyers) place higher bids than theoretically predicted.

than predicted in treatment *B* and 12.2% lower in treatment *S*. This pattern of underbidding continues to hold even when we distinguish by supplier's role: in treatment *B*, the innovator's bids are 17.5% lower than predicted, while those of the challenger are 12.6% lower; in treatment *S*, instead, underbidding is relatively similar between the innovator (11.5%) and the challenger (12.9%).

Table 5 (columns 1 and 2) parametrically investigates the determinants of bids, pooling data from both treatments. Estimates in column 1 show that there is no statistically significant difference in bids between the two treatments, an evidence that further supports the rejection prediction (P2). The specification in column 2 allows us to test prediction (P5) concerning the determinants of bidding behavior in the auction stage of *PING*. First, the effect of the production cost is as expected (prediction P5-a): a higher production cost translates into a higher bid ($p < 0.001$). It is worth noticing that this effect is approximately 10% weaker in treatment *B* than in treatment *S*, as highlighted by the negative and significant value of the coefficient of $B \times c_{it}$ ($p = 0.036$). Second, in line with with predictions (P5-b) and (P5-c), the innovator's bid is significantly higher than the challenger's one in treatment *B* (for $innovator + B \times innovator$, $p < 0.001$), but not in treatment *S* (for $innovator$, $p = 0.400$). Third, comparing bids between treatments, we find support for prediction (P5-d): the bid of the innovator in treatment *B* is significantly higher than in treatment *S* (for $B \times innovator$, $p < 0.001$); and the bid of the challenger in treatment *B* is significantly lower than in treatment *S* (for *B*, $p = 0.036$).

The previous empirical observations are summarized below.

Result 3. Bidding behavior in *PING*. *Bids are increasing in the production cost. The presence of the bid preference makes the innovator place higher bids than the challenger. Bids in treatment S are in between the innovator's and the challenger's bids in treatment B. All subjects tend to underbid relative to what was theoretically predicted.*

5.3 Bids and sunk cost fallacy

5.3.1 A parametric overview

The results of the auction stage in *PING* leave us with two empirical puzzles. First, in both treatments, *B* and *S*, we observe a general tendency to place bids that are lower than the theoretical levels. Second, and more importantly, the winning bid in treatment *B* is not higher than in treatment *S*, which contradicts prediction (P2). This latter puzzle is crucial because it has relevant welfare consequences, as it implies that the bid preference scheme does not imply any extra-cost for the buyer.

While overly aggressive bidding is a largely documented phenomenon in experimental auctions and is mainly attributed to risk aversion, our puzzling result that winning bids are the same in the two treatments calls for other potential (behavioral) explanations. We conjecture that bidders' behavior in the auction stage might have been influenced by the investments made in the preceding innovation stage, although these investments, being sunk, should not affect subsequent decisions. In other words, we suspect that the sunk cost fallacy may play a role.

Table 5. Bids in the auction stage of *PING*: parametric results.

Treatment	*B* & *S*		*S*
	(1)	(2)	(3)
B	1.054	-5.561**	
	(2.027)	(2.656)	
c_{it}		0.706***	0.706***
		(0.021)	(0.0168)
$B \times c_{it}$		-0.061**	
		(0.029)	
$k(i_{it})$			0.517***
			(0.134)
I		1.080	-0.795
		(1.283)	(1.160)
$B \times I$		17.001***	
		(1.785)	
Constant	66.747***	31.816***	30.28***
	(1.433)	(1.867)	(1.861)
Observations	2160	2160	1080

Notes. This table reports parametric estimates (standard errors in parentheses) from two-way linear random-effects models accounting for both potential individual dependency over periods and dependency within rematching group. The dependent variable is b_{it}, the bid place by the bidder in a period of the auction stage. Columns 1 and 2 pool data from both treatments B and S, column 3 considers treatment S only. B is a treatment dummy, c_{it} is the production cost assigned to the bidder in the period, I is a dummy variable that is equal to 1 if the subject is the innovator in the period and 0 otherwise, $B \times I$ is an interaction term, and $k(i_{it})$ is the investment cost borne by the subject in the innovation stage in the period. Significance levels are denoted as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The sunk cost fallacy is a behavioral bias that is invoked whenever the decision of an individual is influenced by the amount of the investment (in money, effort, or time) she has previously sunk in that course of action, despite the fact that a rational decision maker should ignore the initial investment and concentrate solely on the effects of the current decision. This broad definition does not say anything about the direction of this bias. In fact, both directions have been documented in the literature. Most often, the sunk cost fallacy is referred to as an "escalation of commitment", whereby the decision maker tends "to continue an endeavor once an investment has been made, even if it is not optimal to do so" (Arkes and Blumer, 1985). Generally speaking, the individual is more willing to take an action that confirms or justifies the initial investment. This *standard* view of the sunk cost fallacy has also been named *Concorde effect*, inspired by the fact that France and Britain continued to invest in the Concorde supersonic jet, albeit it was known it was going to be unprofitable.[23] However, more recent evidence has also documented a *reverse* sunk cost fallacy, whereby a higher investment made in a certain course of action makes the decision maker *less* willing to continue on that project, as if a large sunk cost reduced the attractiveness of the project with respect to quitting it (the individual does not want "to throw good money after bad"). Interestingly, this reverse sunk cost fallacy could explain full cost pricing, the practice by which firms often incorporate sunk costs in their pricing decisions, as if part of the fixed cost were converted into variable cost (for this reason, this reverse sunk cost fallacy is also named *pro-rata effect*).

In the context of *PING*, the standard version of the sunk cost fallacy would imply that a subject who borne a larger sunk cost in the innovation stage places a lower bid in the auction stage, in order to increase the winning probability: in fact, only winning the auction would psychologically confirm that the initial investment in innovation was not useless. On the other hand, the reverse sunk cost fallacy would manifest itself in a positive relationship between the sunk investment cost and the bid placed in the auction, much like a firm that adopts full cost pricing. Thus, once empirically rejected the null prediction stated in P3, the sunk cost fallacy argument leads us to conjecture the following two alternative hypotheses.[24]

(P3′) STANDARD SUNK COST FALLACY. *The bidding functions in the auction stage of PING are decreasing in the investment made in the innovation stage.*

(P3″) REVERSE SUNK COST FALLACY. *The bidding functions in the auction stage of PING are increasing in the investment made in the innovation stage.*

To test the empirical validity of this behavioral argument, we first analyze *MG*, as this game provides a clean causal test of the presence of the sunk-cost fallacy, which is immune to confounding factors. In fact, like *PING*, *MG* involves an initial investment decision –

[23]Several psychological and cognitive theories have been proposed to explain this behavior. For example, individuals may have a taste for rationalizing past decisions, or they may fear to lose reputation if they reverse course on an initial investment. Baliga and Ely (2011) suggest that, in case of imperfect memory, a large sunk initial investment may suggest that the project was estimated to be particularly valuable and is thus worth continuing.

[24]Hypothesis (P3′) and (P3″) correspond to the pre-registered hypothesis (HP3-3) and (HP3-2), respectively (https://aspredicted.org/dg95u.pdf).

whether or not to enter the market – that generates sunk costs, followed by a bidding choice. However, differently from to *PING*, in *MG*: (i) there is no strategic interaction, as each subject bids against a randomly generated offer by the computer, and (ii) by virtue of the *BDM* mechanism, the theoretical bidding strategy is unaffected by the subject's risk attitude, and the selling offer should coincide with the corresponding production cost.

Parametric results of the basic specifications are presented in the first two columns of Table 6. Two main findings emerge from the table. First, while the seller's production cost γ_{it} assigned in the second stage is a strong and positive determinant of her selling offer, the magnitude of the coefficient is significantly lower than its predicted value of 1 (around 0.6, $p < 0.001$). In other words, while the positive relationship ensures the subjects' understanding of the experimental task, the coefficient's magnitude suggests that, on average, sellers only incorporate about 60% of a change in their cost into their offer. Second, the entry cost e_{it} assigned in the first stage has a significant and positive impact on the selling offer in the subsequent stage ($p = 0.001$).

These findings are robust to several alternative specifications. First, the results are not affected by the treatment condition the subjects were assigned to in *PING* (for the coefficient of S in column 2, $p = 0.301$). Second, results remain virtually unchanged when we estimate a Tobit model to account for the fact that the selling offer in the second stage of *MG* is truncated at 0 (see columns 3 and 4).

Finally, we address a potential sample selection bias in the analysis of the selling offers. All regressions in columns 1-4 use data from the second stage of *MG*, which are available only from those subjects who chose to enter the market and bear the corresponding entry cost. For this reason, in columns 5 and 6, we estimate a Heckman model in which the selection equation shapes the entry decision. The determinants of the entry decision are the assigned entry cost, along with measures of risk tolerance and knowledge of game theory as elicited in the post-experiment questionnaire. The selection variables are all significant and exhibit expected signs: a higher entry cost decreases the probability of entry, while greater risk tolerance increases it; finally, a prior knowledge of game theory reduces the probability of entry. Crucially, after controlling for potential selection bias, the results from the second-stage regression are fully confirmed: the (sunk) entry cost remains a positive and statistically significant determinant of the selling offer. Moreover, the estimated correlation between the error terms of the selection and the outcome equations (ρ) is negative and highly significant ($\rho = -0.394$, $p < 0.001$). This finding confirms the presence of a sample selection bias and suggests that unobserved factors that discourage the subject's entry decision are associated with higher selling offers bids among those who effectively entered the market.

All in all, the parametric investigation presented above suggests that although the entry cost borne in the first stage of *MG* is sunk, it strongly and positively affects the selling offers made by a subject in the second stage, thus providing *prima facie* evidence in favor of the reverse sunk cost fallacy hypothesis (P3").

Result 4. Selling offers and entry (sunk) costs in *MG*. *In MG, selling offers are positively affected by the (sunk) entry costs borne in the previous stage.*

We can now turn to the analysis of bids in the auction stage of *PING* and assess whether

Table 6. Determinants of the selling offers in *MG*: parametric results.

	RE		Tobit		Heckman	
	(1)	(2)	(3)	(4)	(5)	(6)
Stage 2: Selling Offer						
γ_{it}	0.588***	0.588***	0.645***	0.645***	0.586***	0.585***
	(0.033)	(0.033)	(0.037)	(0.037)	(0.033)	(0.033)
e_{it}	0.217***	0.217***	0.229***	0.230***	0.262***	0.261***
	(0.063)	(0.063)	(0.070)	(0.070)	(0.090)	(0.090)
S		-0.941		-0.958		-1.038
		(0.910)		(1.088)		(0.873)
Constant	3.951***	4.425***	2.963***	3.448***	4.023***	4.549***
	(0.556)	(0.720)	(0.653)	(0.853)	(0.564)	(0.716)
Stage 1: Entry Decision (Heckman)						
Risk tolerance					0.190**	0.187**
					(0.074)	(0.074)
Game theory					-0.608***	-0.619***
					(0.224)	(0.224)
e_{it}					-0.573***	-0.573***
					(0.034)	(0.034)
Constant					3.860***	3.886***
					(0.462)	(0.464)
ρ (rho)					-0.394***	-0.400***
					(0.099)	(0.099)
Observations	971	971	971	971	1440	1440

Notes. This table reports estimates from different parametric specifications (standard errors are reported in parentheses) with random effects that account for dependency of subject's choices across periods. Columns (1) and (2) report results from GLS models. Columns (3) and (4) report results from a Tobit model. Columns (5) and (6) report results from a Heckman model. Results from the selection model for the entry decision in the first stage of *MG* are also reported in the second part of the table. In all columns, the dependent variable is the selling offer in the second stage of *MG*. γ_{it} is the seller's production cost assigned in the second stage. e_{it} is the entry cost assigned in the first stage, and S is a dummy equal to 1 in case the subject played treatment S in *PING* in that period. *Risk tolerance* is an ordinal discrete variable capturing the level of risk tolerance as reported by the subject in the post-experimental questionnaire. *Game theory* is a dummy variable that takes a value of 1 if the subject attended lessons of game theory during her university studies. Significance levels are denoted as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

and to what extent the observed behavior confirms the validity of the sunk cost fallacy detected in *MG*. To do this, we should establish a causal relationship between $k(i_{it})$, the investment cost borne in the innovation stage by individual i in period t, and b_{it}, the bid placed by that individual in the auction stage in the same period.

The simplest approach is to estimate, separately for each treatment, the following model:

$$b_{it} = \beta_c c_{it} + \beta_k k(i_{it}) + \beta_I I_{it} + \iota_i + \kappa_i + \epsilon_{it} \tag{1}$$

where c_{it} is the production cost randomly assigned to subject i in the auction stage, I_{it} is a dummy variable indicating whether subject i is the innovator in the period, and ι_i and κ_i are random effects accounting for both individual dependence over repeated auctions and dependency within rematching groups. However, equation (1) cannot be realiably estimated using a linear model because the investment cost, $k(i_{it})$, affects both the bid, b_{it}, through the (possible) sunk cost fallacy, and the probability of being the innovator, I_{it}, through the level of innovation chosen, i_{it}.

This endogeneity concern is substantially less severe in treatment *S*, where the suppliers' roles (innovator or challenger) have no effects on the auction rules and thus should not affect the bidding strategies. We therefore proceed by estimating equation (1) for treatment *S* alone. Results are presented in Column 3, Table 5. Notice first that, as predicted by the theory (prediction P5-c), the innovator status plays no role in determining bids (for the coefficient of I, $p = 0.493$), a finding that alleviates our primary endogeneity concern. Second, as expected (prediction P5-a), the production cost, c_{it}, positively and significantly affects a subject's bid. Crucially, we find that the investment sunk cost k_{it} borne in the innovation stage of *PING* significantly increases bids in the procurement stage (the coefficient of $k(i_i t)$ is positive and significant: $p < 0.001$). In line with what emerged from *MG*, this result parametrically confirms the presence of the reverse sunk cost fallacy (hypothesis P3″). As discussed above, the endogenous nature of k_{it} led us to refrain from extending the parametric analysis to treatment *B*. For this reason, in the next part of the analysis, we rely on a structural approach to identify and isolate the sunk cost fallacy in treatment *B* of *PING*.

5.3.2 Structural analysis: a Quantal Response Equilibrium model

We now turn to a structural approach evidence from a Quantal Response Equilibrium (QRE) model with three parameters. In a QRE model, the assumption that a subject always chooses her best response to the opponents' strategies is replaced by a probabilistic choice function tuned by an error parameter: the probability of playing a suboptimal strategy is strictly positive, but it depends on the (relative) payoff associated with it. In other words, the condition of full rationality is weakened by allowing for errors, but such errors are somewhat rational in the sense that the probability of playing a suboptimal strategy decreases in the level of the payoff associated with it. Importantly, the QRE concept preserves the equilibrium condition that beliefs are correct, meaning that a subject correctly anticipates that opponents are playing such probabilistic strategies involving errors.

Given the dynamic nature of *PING*, to compute the QRE we impose belief consistency and sequential rationality (weakened in the QRE sense) at each information set of the auction stage (see McKelvey and Palfrey, 1998). Specifically, let $p_s(b, c, i_s, \mathbf{i}_{-s}, \mathbf{g})$ denote the probability that a subject of status $s \in \{I, C\}$ bids b, when her production cost is c, she chose the innovation level i_s in the innovation stage, she has information $\mathbf{i}_{-s}$ on the innovation level chosen by the opponent and holds *prior* beliefs $\mathbf{g}$ on the innovation choice made by the opponent (i.e. $\mathbf{g} = \{g_0, g_1, \ldots, g_4\}$, where $g_i \in [0, 1]$ is the prior belief that the opponent chooses the innovation level i in the innovation stage). It is intended that, upon bidding in the auction, these prior beliefs are correctly updated through Bayes' rule using the information $\mathbf{i}_{-s}$. Notice, in particular, that, while the challenger is always informed about the exact investment made by the opponent (i.e. for the challenger, $\mathbf{i}_I$ is a scalar, and thus she knows that the innovator has chosen that level of innovation regardless of the prior beliefs $mathbf g$), the innovator is not informed on the exact investment made by the opponent, but only knows that $i_C \leq i_I$. Hence, in $p_I(b, c, i_I, \mathbf{i}_C, \mathbf{g})$, $\mathbf{i}_C$ is the vector of possible innovation levels no greater than i_I, and, thus, the innovator's posterior beliefs will be a probability distribution over these values, obtained from $\mathbf{g}$ through Bayes' rule.

Now, for each s, b, c, i_s, $\mathbf{i}_{-s}$, and $\mathbf{g}$, $p_s(b, c, i_s, \mathbf{i}_{-s}, \mathbf{g})$ is the solution to the following (logistic) equation:

$$p_s(b, c, i_s, \mathbf{i}_{-s}, \mathbf{g}) = \frac{\exp\{U(s, b, c, i_s, \mathbf{i}_{-s}, \mathbf{g}|\mathbf{p}_{-s})/\mu\}}{\sum_{b \in B_c} \exp\{U(s, b, c, i_s, \mathbf{i}_{-s}, \mathbf{g}|\mathbf{p}_{-s})/\mu\}}, \tag{2}$$

where μ is the QRE error parameter, and $U(s, b, c, i_s, \mathbf{i}_{-s}, \mathbf{g}|\mathbf{p}_{-s})$ is the expected utility of a subject with status s, who bids b in the auction, has production cost c, has chosen innovation level i_s in the innovation stage, has information $\mathbf{i}_{-s}$ on the innovation choice made by the opponent, conditional on the fact that the other subject bids according to the QRE probabilities $\mathbf{p}_{-s}$ as described by equation (2) above. In (2), B_c is the set of admissible and individually rational bids for a subject with production cost c.[25] Clearly, the system (2) has a solution for each possible vector of prior beliefs $\mathbf{g}$. To close the model, we impose that the beliefs must be consistent with actual choices: the vector of prior beliefs $\mathbf{g}$ must coincide with the actual probability distribution of innovation levels in the innovation stage:

$$g_i = \frac{\exp\{U(i|\mathbf{p}_I, \mathbf{p}_C, \mathbf{g})/\mu\}}{\sum_{j=0}^{4} \exp\{U(j|\mathbf{p}_I, \mathbf{p}_C, \mathbf{g})/\mu\}}, \tag{3}$$

where $U(i|\mathbf{p}_I, \mathbf{p}_C, \mathbf{g})$ is the expected utility of a subject who chooses innovation level i, taking into account that, depending on the outcome of the innovation stage, she and the opponent will place bids in the auction according to the probabilities $\mathbf{p}_I$ and $\mathbf{p}_C$ as described by (2), and conditional on the fact that the other subject chooses i according to the QRE probabilities $\mathbf{g}$ as described by equation (3) above.

To incorporate the sunk cost fallacy component in the QRE model in a parsimonious way, we proceed as follows. The intrinsic feature of a sunk cost is that it is borne by the individual *in any case*, whatever her future choices will be, and regardless of the realization of the uncertainty. Hence, despite the actual psychological motives behind this bias, the sunk cost fallacy emerges whenever the subject behaves as if the sunk cost is not really sunk, i.e. exerts its effect differently depending on what happens next. We therefore introduce a sunk cost parameter $\psi \in [0, 2]$, which weighs the investment cost differently

[25] Clearly, for all c, i_s, $\mathbf{i}_{-s}$, $\mathbf{g}$, it must be $\sum_{b \in B_c} p_s(b, c, i_s, \mathbf{i}_{-s}, \mathbf{g}) = 1$.

depending on whether the subject wins or loses the auction. Specifically, the perceived payoff of a subject who made the sunk investment k, whose production cost is c, who bids b and wins the auction is:

$$w_W = b - c - \psi \times k, \tag{4}$$

while, in case she loses the auction, it is:

$$w_L = -(2 - \psi) \times k. \tag{5}$$

In particular, $\psi = 1$ corresponds to the situation in which the subject is not affected by the sunk cost fallacy. The standard sunk cost fallacy (or Concorde effect) manifests itself when $\psi < 1$: the individual perceives a smaller sunk cost when she wins the auction, as if winning the auction would allow her to avoid the investment cost, at least partially. This leads the subject to bid more aggressively in order to secure winning the auction, thereby justifying the initial investment. When, instead, $\psi > 1$, the reverse sunk cost fallacy (or pro-rata effect) arises: the individual perceives a smaller sunk cost when she does not win the auction, as if part of the sunk cost were converted into a variable one. This leads the subject to bid less aggressively to secure a satisfactory payoff in case of winning the auction at the cost of reducing the chance of winning.

Finally, to allow for possible departures from risk neutrality, that may potentially rationalize the underbidding that is observed in both treatments, we consider a Constant Absolute Risk Aversion (CARA) ex-post utility function for suppliers. In particular, the ex-post utility of a supplier is

$$u(w_j) = \frac{1 - \exp\{-r \cdot w_j\}}{r}$$

where r is the risk aversion coefficient and w_j, $j = W, L$, is the perceived payoff dependent on the outcome of the auction, as defined in (4) and (5).[26]

The results of the QRE model fitted to our experimental data are presented in Table 7. The baseline model, which is estimated under the assumption of no sunk cost fallacy ($\psi = 1$), produces similar estimates of the risk aversion parameter in the two treatments, which are greater than zero: subjects tend to be risk averse, a result that is in line with the tendency to aggressively underbid with respect to what theoretically predicted under risk neutrality.

The augmented model allows for the presence of the sunk cost fallacy. The value of the likelihood function associated with the augmented model increases significantly with respect to the baseline model, and the estimates of the risk aversion parameters in the two treatments do not change significantly. On the other hand, for both treatments, the estimated value of ψ is significantly larger than 1 for both treatments, which constitutes evidence in favor of the reverse sunk cost fallacy: in *PING*, subjects' bids in the auction stage are increasing in the (sunk) investment cost borne in the innovation stage.

Interestingly, the estimated value of ψ is significantly larger in treatment S than in treatment B. This means that, while in both treatments the reverse sunk cost fallacy induces the

[26]Higher values of r indicate a higher degree of (absolute) risk aversion. When $r \to 0$, $u(w) \to w$, namely the agent is risk neutral.

Table 7. Baseline and augmented QRE models: Estimated parameters.

Model	Treatment	r	μ	ψ	Log-Likelihood
Baseline	S	0.10	1.25	-	-1545.371
Baseline	B	0.15	1.35	-	-1268.353
Augmented	S	0.05	0.37	1.9	-1313.42
Augmented	B	0.15	1.1	1.5	-1182.954

Notes. This table reports the estimated parameters of both the baseline and the augmented QRE models. The augmented QRE model incorporates the sunk cost fallacy parameter $\psi \in [0, 2]$. r is the CARA risk aversion parameter, and μ is the error parameter. Precision (of the variables): $r \pm 0.05$ (0.01 checked as well), μ model/treat. dependent, $\psi \pm 0.1$

suppliers to bid higher after a larger investment in innovation, this upward adjustment is stronger in treatment S. This fact can justify why, the expected difference between the winning bid in treatment B and in treatment S (prediction (P2)) is not observed in the data.

One possible explanation for this differential intensity of the reverse sunk cost fallacy may be related to the different level of complexity between the two treatments. In treatment B, bids are treated asymmetrically to determine the auction winner, which increases the subjects' cognitive load. This additional complexity captures the attention and the cognitive resources of the participants, reducing the salience of the investment cost just borne and thus weakening their response to the sunk cost fallacy. In contrast, treatment S is a standard first-price auction: the cognitive load is lighter, the sunk cost becomes more prominent, and the sunk cost fallacy emerges more strongly.

Together, the parametric results presented in the previous section and estimates from the QRE model are summarized in the following statement.

Result 5. Bids and innovation (sunk) costs in *PING*. *In PING, bids in the auction stage are positively affected by the (sunk) investment costs borne in the innovation stage. This effect is stronger when the bid preference is absent.*

6 Conclusion

Scholars and practitioners have recognized the potential of public procurement as a driver of innovation. This paper studies a procurement practice meant to stimulate innovation without sacrificing competition in the awarding phase: granting a bid preference in the auction to the firm that submitted the most innovative proposal in the first place. The theory shows that this scheme effectively boosts innovation, but at the cost of a higher awarding price in the auction and of some inefficiency in the allocation.

This paper explores this trade-off through an experiment that mimics such an incentive scheme. Our interest was mainly motivated by the fact that bidders compete for the awarding of the contract after they have sunk the investment costs required to come out with an innovative proposal. Hence, we suspect that the well documented sunk cost fallacy might somehow affect the outcome of this incentive scheme.

The results of our experiment confirm that the introduction of a bid preference scheme generates more innovation; however, in contrast to the predictions, this does not come at a higher awarding price. We show that this puzzling evidence can be explained by the presence of a *reverse sunk cost fallacy*, which operates with different strengths when the bid preference is present and when it is absent. Thus, our work suggests that the negative side of the trade-off associated with bid preference schemes seems to be less worrying than expected. In this respect, our results provide formal support to what has been documented in the sparse cases where these incentive schemes have been adopted in real world contexts. For example, in public procurement for smart city projects, such as those seen in Barcelona and Singapore, innovation-driven bid preferences have accelerated the adoption of advanced technologies without significantly increasing costs (Carboni and Russel, 2019; OECD, 2017). Similarly, in the healthcare sectors, where initial investments are high, such schemes have promoted technological breakthroughs while maintaining competitive pressure (Snowdon et al., 2019; EcoQUIP Plus Consortium, 2024).

More generally, our analysis suggests that any policy meant to promote innovation should take into account that, since innovation involves large investments made upfront, the sunk cost fallacy – a bias documented not only in the lab but also in real world settings – could significantly affect the effectiveness and, ultimately, the success of the policy itself.

References

Arkes, H. R. and P. Ayton (1999). The sunk cost and concorde effects: Are humans less rational than lower animals? *Psychological bulletin 125*(5), 591.

Arkes, H. R. and C. Blumer (1985). The psychology of sunk cost. *Organizational behavior and human decision processes 35*(1), 124–140.

Arrow, K. J. (1972). *Economic welfare and the allocation of resources for invention*. Springer.

Athey, S., D. Coey, and J. Levin (2013). Set-asides and subsidies in auctions. *American economic journal: microeconomics 5*(1), 1–27.

Augenblick, N. (2016). The sunk-cost fallacy in penny auctions. *The Review of Economic Studies 83*(1), 58–86.

Baliga, S. and J. C. Ely (2011). Mnemonomics: the sunk cost fallacy as a memory kludge. *American Economic Journal: Microeconomics 3*(4), 35–67.

Becker, G. M., M. H. DeGroot, and J. Marschak (1964). Measuring utility by a single-response sequential method. *Behavioral science 9*(3), 226–232.

Bennett, J. and E. Iossa (2006). Building and managing facilities for public services. *Journal of public economics 90*(10-11), 2143–2160.

Bhattacharya, V. (2021). An empirical model of r&d procurement contests: An analysis of the dod sbir program. *Econometrica 89*(5), 2189–2224.

Cabral, L., G. Cozzi, V. Denicolo, G. Spagnolo, and M. Zanza (2006). Procuring innovation.

Carboni, J. and D. Russel (2019). Smart city procurement: Innovation, sustainability, and policy challenges. *Journal of Public Procurement 19*(4), 313–331.

Che, Y.-K., E. Iossa, and P. Rey (2021). Prizes versus contracts as incentives for innovation. *The Review of Economic Studies 88*(5), 2149–2178.

Chen, B. R. and Y. S. Chiu (2010). Public–private partnerships: task interdependence and contractibility. *International Journal of Industrial Organization 28*(6), 591–603.

Chiappinelli, O., L. M. Giuffrida, and G. Spagnolo (2025). Public procurement as an innovation policy: Where do we stand? *International Journal of Industrial Organization*, 103157.

Chiappinelli, O. and G. Seres (2024). Optimal discounts in green public procurement. *Economics letters 238*, 111705.

Corns, A. and A. Schotter (1999). Can affirmative action be cost effective? an experimental examination of price-preference auctions. *American Economic Review 89*(1), 291–305.

Decarolis, F., G. de Rassenfosse, L. M. Giuffrida, E. Iossa, V. Mollisi, E. Raiteri, and G. Spagnolo (2021). Buyers' role in innovation procurement: Evidence from us military r&d contracts. *Journal of Economics & Management Strategy 30*(4), 697–720.

EcoQUIP Plus Consortium (2024). Delivering efficiency, quality and sustainability in healthcare through innovation procurement: Case study report. Technical report, European Commission, Public Buyers Community.

Eswaran, M. and H. M. Neary (2016). The evolutionary logic of honoring sunk costs. *Economic Inquiry 54*(2), 835–846.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics 10*(2), 171–178.

Friedman, D., K. Pommerenke, R. Lukose, G. Milam, and B. A. Huberman (2007). Searching for the sunk cost fallacy. *Experimental Economics 10*(1), 79–104.

Griffith, R., S. Redding, and J. V. Reenen (2004). Mapping the two faces of r&d: Productivity growth in a panel of oecd industries. *Review of economics and statistics 86*(4), 883–895.

Hartig, B. (2017). Two wrongs make a right?! is the sunk-cost-effect a commitment device against present bias? *Is the Sunk-Cost-Effect a Commitment Device Against Present Bias*.

Ho, T.-H., I. P. Png, and S. Reza (2018). Sunk cost fallacy in driving the world's costliest cars. *Management Science 64*(4), 1761–1778.

Hodges, J. T. and G. Dellacha (2007). Unsolicited infrastructure proposals: How some countries introduce competition and transparency.

Hong, F., W. Huang, and X. Zhao (2019). Sunk cost as a self-management device. *Management Science 65*(5), 2216–2230.

Hoppe, E. I., D. J. Kusterer, and P. W. Schmitz (2013). Public–private partnerships versus traditional procurement: An experimental investigation. *Journal of Economic Behavior & Organization 89*, 145–166.

Hoppe, E. I. and P. W. Schmitz (2013). Public-private partnerships versus traditional procurement: Innovation incentives and information gathering. *The RAND Journal of Economics 44*(1), 56–74.

Howell, S. T. (2017). Financing innovation: Evidence from r&d grants. *American economic review 107*(4), 1136–1164.

Iossa, E., F. Biagi, and P. Valbonesi (2018). Pre-commercial procurement, procurement of innovative solutions and innovation partnerships in the eu: rationale and strategy. *Economics of Innovation and New Technology 27*(8), 730–749.

Iossa, E. and D. Martimort (2012). Risk allocation and the costs and benefits of public–private partnerships. *The RAND Journal of Economics 43*(3), 442–474.

Krasnokutskaya, E. and K. Seim (2011). Bid preference programs and participation in highway procurement auctions. *American Economic Review 101*(6), 2653–86.

Marion, J. (2007). Are bid preferences benign? the effect of small business subsidies in highway procurement auctions. *Journal of Public Economics 91*(7-8), 1591–1624.

McAfee, R. P. and J. McMillan (1989). Government procurement and international trade. *Journal of International Economics 26*(3), 291–308.

McKelvey, R. D. and T. R. Palfrey (1998). Quantal response equilibria for extensive form games. *Experimental economics 1*(1), 9–41.

Molden, D. C. and C. M. Hui (2011). Promoting de-escalation of commitment: A regulatory-focus perspective on sunk costs. *Psychological Science 22*(1), 8–12.

Negrini, M., A. Riedl, and M. Wibral (2022). Sunk cost in investment decisions. *Journal of Economic Behavior & Organization 200*, 1105–1135.

OECD (2017). Public procurement for innovation – good practices and strategies. *OECD Publishing, Paris*.

Phillips, O. R., R. C. Battalio, and C. A. Kogut (1991). Sunk and opportunity costs in valuation and bidding. *Southern Economic Journal*, 112–128.

Slavtchev, V. and S. Wiederhold (2016). Does the technological content of government demand matter for private r&d? evidence from us states. *American Economic Journal: Macroeconomics 8*(2), 45–84.

Snowdon, A. W., R. Axler, M. St. Pierre, and R. DeForge (2019). Innovation procurement in health systems: Exploring practice and lessons learned. *Healthcare Quarterly 22*(3), 10–13.

Staw, B. M. (1976). Knee-deep in the big muddy: A study of escalating commitment to a chosen course of action. *Organizational behavior and human performance 16*(1), 27–44.

Strough, J., C. M. Mehta, J. P. McFall, and K. L. Schuller (2008). Are older adults less subject to the sunk-cost fallacy than younger adults? *Psychological Science 19*(7), 650–652.

Whyte, G. (1986). Escalating commitment to a course of action: A reinterpretation. *Academy of management review 11*(2), 311–321.

A Instructions for treatment *B* of *PING*

The original instructions were written in Italian. The instructions for *PING* were distributed at the beginning of the session, while those for *MG* were handed out after the final period of *PING*. In the instructions, *PING* and *MG* are referred to as 'Experiment 1' and 'Experiment 2', respectively.

INSTRUCTIONS

- Welcome. Thank you for participating in this experimental session. By following the instructions carefully, you will be able to earn an amount in euros that will be paid to you in cash at the end of the session.

- 24 subjects participate in this session. During the session, it is not allowed to communicate in any way with other subjects. If you have any questions, raise your hand and one of the assistants will come to answer it. The following rules are the same for everyone.

- In this session, you will participate in two distinct experiments. The following instructions are for the first experiment. The instructions for the second experiment will be given to you at the end of the first experiment.

- Although you will participate in two experiments, the payment you will get at the end of the experimental session will be determined based on the earnings you earn in only one of the two experiments. In particular, the experiment used to determine your final payment will be drawn randomly by flipping a coin at the end of the session. In addition, you will be awarded 4 euros for showing up on time for the session.

EXPERIMENT #1

GENERAL RULES

- The first experiment includes 15 periods. At the beginning of each period you will be randomly and anonymously assigned to a group consisting of you and another participant. The composition of your group will change in every period so that you will never interact with the same participant for two consecutive periods.
- Both the identity and the final payments of the subjects with whom you interact during the experiment will remain anonymous. In addition, during the experiment, you and the other participant in the group will make your choices simultaneously.
- In each period, you will be given an initial endowment of 16 points. Your earnings over the period will be given by the initial allocation of 16 points plus the additional points you earned during that period.
- Your total earnings from the first experiment will be the sum of the points you earned over the 15 periods. At the end of the experiment, if the first experiment is selected to determine the final payments, the points you earned will be converted into euros at the exchange rate of 33 points = 1 euro.

THE ECONOMIC SITUATION

- Two sellers compete in an auction to sell an item to a hypothetical buyer. Both you and the other participant assigned to your group are the sellers. On the contrary, no participant plays the role of buyer. The interaction between the two sellers is developed in two consecutive steps.
- In STEP 1, each seller makes an investment choice. The investment is costly and each seller incurs a cost that increases as the level of investment chosen increases. The sellers' investment choices have two consequences on the rules of the auction in STEP 2:
  - the seller who made the highest investment choice plays the role of INVESTOR and will enjoy an advantage in the auction conducted in STEP 2;
  - the investment chosen by the INVESTOR also determines the probability of the auction being revoked in STEP 2: the higher the investment chosen by the INVESTOR, the lower the probability that the auction will be revoked.
- In STEP 2, the two sellers participate in the auction to sell the item and choose the prices at which they are willing to sell the item to the buyer. The outcome of the auction is determined by comparing the price chosen by the INVESTOR with the price of the other seller increased by 50%: the lower of these two values determines the winner of the auction. The winner of the auction sells the item to the buyer and bears the cost of transferring the item. At the end of the period, based on the choice made by the INVESTOR in STEP 1, the computer determines whether or not

to revoke the auction: in case of revocation, the earnings obtained from participation in the auction in STEP 2 are cancelled.

- The rules of STEP 1 and STEP 2 are described in detail below.

STEP 1 - INVESTMENT CHOICE

- In STEP 1, you and the other seller in the group will make your investment choices. Specifically, each seller will make his/her investment choice by choosing an integer number between 0 (lowest investment) and 4 (highest investment).
- The computer will then compare your investment choice with the choice made by the other seller in the group. The seller who has made the highest investment choice will play the role of INVESTOR, while the other will play the role of NON-INVESTOR. In the event of a tie, the computer will assign the two roles randomly and with equal probability.
- Your investment choice is expensive. In particular, the cost you will incur based on your investment choice will be calculated as follows:

  $$\textit{Cost of investment} = (\textit{Investment choice})^2$$

  Therefore, the cost of the investment increases more than proportionally to the level of investment you choose and will be given by a number ranging from 0 (when you choose an investment equal to 0) to 16 (when you choose an investment equal to 4).
- At the end of STEP 1, the computer will show you the role you have been assigned, the cost of your investment choice and, if you are the NON-INVESTOR, the investment choice of the seller who plays the role of INVESTOR.

STEP 2 - PRICE CHOICE

- In STEP 2, you and the other seller in the group will participate in the auction to sell the item to the hypothetical buyer. In particular, each seller in the group will choose the price at which he/she is willing to sell the item to the buyer, which must be an integer number between 0 and 135.
- The seller who plays the role of INVESTOR will enjoy an advantage in the auction. In fact, in order to determine the winner of the auction, the computer will compare two values: the price chosen by the INVESTOR with the price chosen by the NON-INVESTOR increased by 50%. The winner of the auction is the seller with the lowest value. In the event of a tie, the computer will randomly select the winner of the auction with equal probability.
- The earnings obtained in STEP 2 by the seller who was not the winner of the auction are null. On the contrary, the earnings obtained in STEP 2 by the winner of the auction are given by:

  $$\textit{Earnings of the winner of the auction} = \textit{Sale price} - \textit{Cost for transferring the item}$$

- The computer will tell you the value of your cost for transferring the item before you choose your selling price. The cost for transferring the item is selected by the computer randomly and with equal probability by drawing an integer between 0 and 100. This means that your cost for transferring the item does not depend on the other seller's cost or on the cost assigned to you in previous periods. In addition, each seller will only be told the value of h cis/her cost for transferring the item, and will not receive any information about the cost assigned to the other seller.

- There is a possibility that the auction will be revoked at the end of the auction. The probability that the auction will be revoked depends on the level of investment chosen by the INVESTOR in STEP 1: the higher the investment chosen by the INVESTOR, the lower the probability that the auction will be revoked. In particular, the probability of revocation is calculated on the basis of the following table:

Investment chosen by the INVESTOR in STEP 1	0	1	2	3	4
Probability of revocation	80%	60%	40%	20%	0%

- The revocation of the auction will be determined on the basis of a lottery. Specifically, at the end of the auction, your group will be assigned a number of tickets ranging from 1 to 80 if the probability of cancellation is 80%, from 1 to 60 if the probability of cancellation is 60%, from 1 to 40 if the probability of cancellation is 40%, from 1 to 20 if the probability of cancellation is 20%, and no tickets if the probability of cancellation is equal to 0%. Next, the computer will randomly and with equal probability draw an integer between 1 and 100. If the ticket drawn is among those allocated to your group, the STEP 2 auction will be revoked.

- In the event of revocation, the earnings obtained by the auction winner are voided.

- The following table summarizes how your earnings in a period will be determined on the basis of the choices you made in STEP 1 and STEP 2.

IF YOU WIN THE AUCTION AND THE AUCTION IS NOT REVOKED	**IF YOU DO NOT WIN THE AUCTION AND THE AUCTION IS NOT REVOKED**	**IF THE AUCTION IS REVOKED**
Earnings in a period = Initial allocation of 16 points − Cost of investment (STEP 1) + Selling price (STEP 2) − Cost for transferring the item (STEP 2)	Earnings in a period = Initial allocation of 16 points − Cost of investment (STEP 1)	Earnings in a period = Initial allocation of 16 points − Cost of investment (STEP 1)

- At the end of the period, the computer will show you the investment choice made by the INVESTOR, the cost you incurred for your investment choice in STEP 1, the price you chose, the outcome of the auction with the winning price in STEP 2, the outcome of the auction revocation lottery, and your earnings for the period expressed in points.

TO SUM UP...

1. During each period of the first experiment, you will participate in two consecutive steps.
2. In STEP 1, you and the other seller in your group will simultaneously make your investment choices, indicating an integer number between 0 (lowest investment) and 4 (highest investment). The seller who has made the highest investment choice will play the role of INVESTOR, while the other seller will play the role of NON-INVESTOR.
3. Your investment choice is costly. In particular, you will incur a cost for your investment choice calculated on the basis of the following expression:

$$\textit{Cost of investment} = (\textit{Investment choice})^2$$

4. In STEP 2 you will participate in the auction for the sale of the item to the hypothetical buyer. Specifically, you and the other seller in the group will simultaneously choose your selling prices, indicating an integer number between 0 and 135.
5. The seller who plays the role of INVESTOR will enjoy an advantage in the auction. In fact, in order to determine the winner of the auction, the INVESTOR's price will be compared with the NON-INVESTOR's price increased by 50%. The lower of these two values will determine the auction winner . The earnings of the seller who was not the winner of the auction are null. On the contrary, the earnings obtained in STEP 2 by the winner of the auction are given by:

$$\textit{Earnings of the winner of the auction} = \textit{Sale price} - \textit{Cost for transferring the item}$$

6. The cost for transferring the item is selected randomly and with equal probability by the computer by drawing an integer between 0 and 100. The information on the cost for transferring the item is private, specific to each seller, and its value changes from period to period.
7. There is a possibility that the STEP 2 auction will be revoked. The higher the level of investment chosen by the INVESTOR in STEP 1, the less likely the auction will be revoked. In the event of a revocation of the auction, the earnings obtained by the sellers in STEP 2 are voided. The following table shows, for each level of investment of the INVESTOR in STEP 1, the probability that the auction will be revoked:

Investment chosen by the INVESTOR in STEP 1	0	1	2	3	4
Probability of revocation	80%	60%	40%	20%	0%

8. The points earned during the period will be determined based on the choices you made in STEP 1 and STEP 2. The following table summarizes how your earnings in a period will be determined.

IF YOU WIN THE AUCTION AND THE AUCTION IS NOT REVOKED	**IF YOU DO NOT WIN THE AUCTION AND THE AUCTION IS NOT REVOKED**	**IF THE AUCTION IS REVOKED**
Earnings in a period = Initial allocation of 16 points − Cost of investment (STEP 1) + Selling Price (STEP 2) − Cost for transferring the item (STEP 2)	Earnings in a period = Initial allocation of 16 points − Cost of investment (STEP 1)	Earnings in a period = Initial allocation of 16 points − Cost of investment (STEP 1)

EXPERIMENT #2

GENERAL RULES

- The second experiment includes 10 periods. Unlike the first experiment, in this second experiment you will not interact with any other participant, but only with the computer.
- In each period, you will be awarded an initial endowment of 9 points. Your earnings over the period will be given by the initial endowment of 9 points plus the additional points you earned during that period.
- Your total earnings from the second experiment will be the sum of the points you earned over the 10 periods. At the end of the experiment, if the second experiment is selected to determine the final payments, the points you earned will be converted into euros at the exchange rate of 12 points = 1 euro.

THE ECONOMIC SITUATION

- You are a seller who has to choose whether to participate in a market where you are given the opportunity to sell an item to a hypothetical buyer. The interaction with the buyer is developed in two consecutive steps.
- In STEP 1, you choose whether or not to participate in the market. Participating in the market is costly and you incur a participation cost. Only if you choose to participate in the market you can access STEP 2, where you will have the opportunity to sell the item to the buyer.
- In STEP 2, you make a selling offer for the item. The outcome of trade is determined by comparing your selling offer with the buying offer made by the buyer. Specifically, if your selling offer is less than or equal to the buyer's buying offer, trade takes place. In this case, you sell the item to the buyer at a price equal to the buying offer and bear the cost for transferring the item. If, instead, your selling offer is greater than the buyer's buying offer, trade does not take place. In this case, you do not sell the item and do not bear any cost for transferring the item.
- The rules of STEP 1 and STEP 2 are described in detail below.

STEP 1 - PARTICIPATION CHOICE

- In STEP 1, you will decide whether or not to participate in the market.
- Before you make this decision, you will be informed of the participation cost. This cost is determined by the computer randomly and with equal probability in each period by drawing an integer number between 0 and 9. The participation cost selected in a period does not depend on the costs drawn in previous periods.

- If you choose not to participate in the market, you will not incur any cost and you will not have to make any other decision. Your earnings in the period will therefore be equal to your initial endowment of 9 points.
- If you choose to participate in the market, you will incur the participation cost and proceed to STEP 2.

STEP 2 - SELLING OFFER

- In STEP 2, you will make your selling offer to sell the item to the buyer.
- The computer will tell you the value of your cost for transferring the item before you choose your selling offer. The cost for transferring the item is the cost you will incur if you sell the item to the buyer. This cost is selected by the computer randomly and with equal probability in each period by drawing an integer between 0 and 15. The cost for transferring the item selected in a period does not depend on the costs assigned to you in previous periods.
- Given the cost for transferring the item, but before knowing the buyer's buying offer, you will choose your selling offer, which must be an integer number between 0 and 30.
- After you have made your selling offer, you will be informed of the buyer's buying offer. This buying offer is selected by the computer in each period randomly and with equal probability by drawing an integer between 0 and 30. The buying offer selected in a period does not depend on those drawn in previous periods, nor on your selling offer.
- The outcome of trade is determined by comparing your selling offer with the buying offer made by the buyer. If your selling offer is greater than the buyer's buying offer, trade does not take place. In this case, you will not sell the item and you will not bear any cost for transferring the item. If, however, your selling offer is less than or equal to the buyer's buying offer, trade takes place. In this case, you will sell the item to the buyer at a price equal to the buying offer and you will bear the cost for transferring the item.
- The following table summarizes how your earnings in a the period will be determined on the basis of the choice you made in STEP 1 and the outcome of the market in STEP 2.

IF YOU CHOOSE NOT TO PARTICIPATE IN THE MARKET	IF YOU CHOOSE TO PARTICIPATE IN THE MARKET AND TRADE DOES NOT TAKE PLACE	IF YOU CHOOSE TO PARTICIPATE IN THE MARKET AND TRADE TAKES PLACE
Earnings in a period = Initial allocation of 9 points	Earnings in a period = Initial allocation of 9 points − Participation cost (STEP 1)	Earnings in a period = Initial allocation of 9 points − Participation cost (STEP 1) + Buyer's buying offer (STEP 2) − Cost for transferring the item (STEP 2)

- At the end of the period, the computer will show you the participation cost and, if you chose to participate, your selling offer, the buyer's buying offer, the cost for transferring the item, and your earnings in that period expressed in points.

TO SUM UP...

1. During each period of the second experiment, you will participate in two consecutive steps.
2. In STEP 1, you will choose whether or not to participate in the market. Before you make your choice, you will be informed of the participation cost, which is selected randomly and with equal probability by the computer by drawing an integer between 0 and 9. If you choose not to participate in the market, you will not incur any cost and will not have to make any other decision. If you choose to participate, you will incur the participation cost and access STEP 2.
3. In STEP 2, you will make your selling offer, indicating an integer number between 0 and 30. Before making your selling offer, you will be informed of the cost for transferring the item. This cost is selected randomly and with equal probability by the computer by drawing an integer between 0 and 15.
4. The outcome of trade depends on the comparison between your selling offer and the buyer's buying offer. The buyer's buying offer is selected randomly and with equal probability by the computer by drawing an integer between 0 and 30. You will be informed of the buyer's buying offer after you have made your selling offer.
5. If your selling offer is greater than the buyer's buying offer, trade does not take place. In this case, you do not sell the item and do not bear any cost for transferring the item.

6. If, however, your selling offer is less than or equal to the buyer's buying offer, trade takes place. In this case, you sell the item to the buyer at a price equal to the buying offer and bear the cost for transferring the item.

7. The points earned during the period will be determined on the basis of the choices you make in STEP 1 and STEP 2. The following table summarizes how your earnings in a period will be determined.

IF YOU CHOOSE NOT TO PARTICIPATE IN THE MARKET	**IF YOU CHOOSE TO PARTICIPATE IN THE MARKET AND TRADE DOES NOT TAKE PLACE**	**IF YOU CHOOSE TO PARTICIPATE IN THE MARKET AND TRADE TAKES PLACE**
Earnings in a period = Initial allocation of 9 points	Earnings in a period = Initial allocation of 9 points − Participation cost (STEP 1)	Earnings in a period = Initial allocation of 9 points − Participation cost (STEP 1) + Buyer's buying offer (STEP 2) − Cost for transferring the item (STEP 2)

B Theoretical predictions: proofs

B.1 Auction stage

TREATMENT S. This is nothing but a symmetric independent private value first-price auction. Using standard techniques, we obtain the following equilibrium bidding function (common to both the Innovator and the Challenger):

$$b(c) = \frac{c + \bar{c}}{2}.$$

TREATMENT B. This is an asymmetric auction in that, although types (i.e. costs) are symmetrically distributed across bidders, their bids are treated asymmetrically. In particular, to establish the winner, the Challenger's bid is augmented by a factor θ (equal to 1.5 in the experiment), but, in case of winning, the Challenger is paid her genuine bid. One could equivalently think that the Challenger submits a bid b_C that is compared with the Innovator's bid, and, in case of winning, the Challenger is paid only b_C/θ. We prefer to use this latter formulation. Now, let $b_I(c_I)$ and $b_C(c_C)$ denote the pure strategy equilibrium bidding functions of the Innovator and the Challenger, respectively. Standard arguments imply that, in the domains of types that have a strictly positive probability of winning the auction, $b_I(c_I)$ and $b_C(c_C)$ are continuous and strictly increasing.

Now, for $j = I, C$, let $[\underline{b}_j, \overline{b}_j]$ be the range of $b_j(\cdot)$, and let $[\underline{b}_j^*, \overline{b}_j^*]$ be (the closure of) the subset of $[\underline{b}_j, \overline{b}_j]$ such that bidder j has a strictly positive probability of winning the auction. (That these sets are indeed intervals is guaranteed by the continuity of $b_I(c_I)$ and $b_C(c_C)$.) Observe that $\underline{b}_I^* = \underline{b}_C^* \equiv \underline{b}^*$ and $\overline{b}_I^* = \overline{b}_C^* \equiv \overline{b}^*$. In fact, if $\underline{b}_I^* < \underline{b}_C^*$, then the Innovator would profit from slightly increasing her bid; and vice versa. On the other hand, $\overline{b}_I^* = \overline{b}_C^*$ follows from the continuity of $b_I(c_I)$ and $b_C(c_C)$.

To determine the equilibrium bidding functions in the region where both bidders have a strictly probability of winning, it is useful to introduce the following functions: let $y_I(\cdot)$ and $y_C(\cdot)$ be the inverses of $b_I(\cdot)$ and $b_C(\cdot)$, respectively (hence, $y_j(b_j)$ is the type of bidder j that, in equilibrium, bids b). Notice that $y_I(\cdot)$ and $y_C(\cdot)$ are both defined on the same domain $[\underline{b}^*, \overline{b}^*]$. Now, the Innovator solves the following optimization problem:

$$\max_b \,(b - c_I)\left[1 - G\left(y_C(b)\right)\right],$$

where $G(\cdot)$ is the probability distribution of the production cost. Similarly, the Challenger chooses her bid to maximize:

$$\max_b \left(\frac{b}{\theta} - c_C\right)\left[1 - G\left(y_I\left(b\right)\right)\right].$$

The two optimization problems above produce the following system of two differential equations:[27]

$$\bar{c} - y_C(b) = \left[b - y_I(b)\right] y_C'(b), \tag{B.1}$$

$$\bar{c} - y_I(b) = \left[b - \theta y_C(b)\right] y_I'(b). \tag{B.2}$$

To solve the system above, we have to pin down the appropriate boundary conditions. To this end, observe that:

- $\underline{b}^* = \underline{b}_j = b_j(\underline{c})$, $j = I, C$ (this follows from the strict monotonicity of $b_j(\cdot)$);

- $\overline{b}^* = \overline{b}_I = b_I(\bar{c}) \in (\bar{c}, \theta\bar{c})$. To see this, consider an Innovator with cost $\bar{c}$: by bidding exactly $\bar{c}$, she would win with strictly positive probability, as all Challenger's types with cost greater than $\bar{c}/\theta$ will bid higher (otherwise they would suffer a loss). Hence, by bidding (slightly) more than $\bar{c}$, the Innovator with cost $\bar{c}$ can obtain a strictly positive expected payoff. On the other hand, by bidding $\theta\bar{c}$ or more, she would win with probability equal to zero, as all Challenger's types below $\bar{c}$ can afford to make a bid slightly below $\theta\bar{c}$, and doing so they would guarantee a strictly positive probability of winning;

- the previous observation implies that there will be a set of Challenger's types – those above $\overline{b}^*/\theta$ – that, in equilibrium, have a zero probability of winning the auction (they cannot afford to bid less than $\overline{b}^*$): we assume these types bid $b_C(c_C) = \theta c_C$;

[27] See equations (4a) and (4b) in Corns and Schotter (1999).

- $y_C(\overline{b}^*) = \overline{b}^*/\theta$. Suppose, instead, that $y_C(\overline{b}^*) < \overline{b}^*/\theta$: take $0 < \varepsilon < (\overline{b}^*/\theta - y_C(\overline{b}^*))/2$. Then type $y_C(\overline{b}^*) + \varepsilon$ could profitably bid $\overline{b}^* - \theta\varepsilon$: her probability of winning will be strictly positive (because her bid is strictly lower than $\overline{b}^*$) and her payoff in case of winning would be strictly positive;

- $\overline{b}^* = (1+\theta)\overline{c}/2$. To see this, consider the problem faced by an Innovator with cost $\overline{c}$: if she bids b (and we know this type will bid strictly above $\overline{c}$), she will defeat all Challenger's types whose cost is higher than b/θ (they cannot afford to bid less than b), while she will lose against all other types (they will find it profitable to bid just below b). Hence, she solves the following maximization problem:

$$\max_{b>\overline{c}} (b-\overline{c})\left[1-G\left(\frac{b}{\theta}\right)\right],$$

which yields $b = (1+\theta)\overline{c}/2$.

- $\underline{b}^* = (1+\theta)\overline{c}/4$. To see this, write equations (B.1) and (B.2) as follows:

$$by_C'(b) + y_C(b) - y_I(b)y_C'(b) = \overline{c},$$
$$by_I'(b) + y_I(b) - \theta y_C(b)y_I'(b) = \overline{c}.$$

Multiply the first equation above by θ and then sum the two equations to obtain

$$\theta\left[by_C'(b) + y_C(b)\right] + \left[by_I'(b) + y_I(b)\right] - \theta\left[y_I(b)y_C'(b) + y_C(b)y_I'(b)\right] = (1+\theta)\overline{c}.$$

Notice that the equation above is equivalent to

$$\theta\left[b\,y_C(b)\right]' + \left[b\,y_I(b)\right]' - \theta\left[y_I(b)\,y_C(b)\right]' = \left[(1+\theta)\,\overline{c}\,b\right]'.$$

Integrating both sides, one obtains

$$\theta\,b\,y_C(b) + b\,y_I(b) - \theta\,y_I(b)\,y_C(b) = (1+\theta)\,\overline{c}\,b + K, \tag{B.3}$$

where K is a constant of integration. To pin down the value of K, we substitute into (B.3) the boundary conditions $y_I(\overline{b}^*) = \overline{c}$ and $y_C(\overline{b}^*) = \overline{b}^*/\theta$, with $\overline{b}^* = (1+\theta)\overline{c}/2$. It obtains $K = -\left[(1+\theta)^2\overline{c}^2\right]/4$. Finally, substituting into (B.3) the conditions $y_I(\underline{b}^*) = y_C(\underline{b}^*) = \underline{c}$, we obtain

$$\underline{b}^* = \frac{(1+\theta)^2\,\overline{c}^2 - 4\theta\underline{c}^2}{4\,(1+\theta)\,(\overline{c}-\underline{c})} = \frac{(1+\theta)\,\overline{c}}{4},$$

where the last equality follows from $\underline{c} = 0$.

Now, the system (B.1)-(B.2), together with the boundary conditions

- $y_I(\underline{b}^*) = y_C(\underline{b}^*) = \underline{c}$, with $\underline{b}^* = (1+\theta)\overline{c}/4$,

- $y_I(\overline{b}^*) = \overline{c}$ and $y_C(\overline{b}^*) = \overline{b}^*/\theta$, with $\overline{b}^* = (1+\theta)\overline{c}/2$,

determines the equilibrium bidding functions (in the region $[\underline{b}^*, \overline{b}^*]$). The solution to this system of differential equations, in terms of the direct functions $b_I(c_I)$ and $b_C(c_C)$, was obtained through numerical techniques, and is depicted in Figure 1.

B.2 Innovation stage

Tables (A1) and (A2) report, for the two treatments, the payoffs of supplier 1 (row player) and supplier 2 (column player) for each possible combinations of innovation choices (i_1, i_2), assuming that suppliers play their equilibrium bids in the Auction stage.

Table A1. Innovation stage: payoff table for treatment S.

	0	1	2	3	4
0	19.3; 19.3	22.7; 21.7	26.0; 22.0	29.4; 20.4	32.7; 16.7
1	21.7; 22.7	21.7; 21.7	25.0; 22.0	28.4; 20.4	31.7; 16.7
2	22.0; 26.0	22.0; 25.0	22.0; 22.0	25.4; 20.4	28.7; 16.7
3	20.4; 29.4	20.4; 28.4	20.4; 25.4	20.4; 20.4	23.7; 16.7
4	16.7; 32.7	16.7; 31.7	16.7; 28.7	16.7; 23.7	16.7; 16.7

Table A2. Innovation stage: payoff table for treatment B.

	0	1	2	3	4
0	19.8; 19.8	20.5; 25.8	22.7; 28.1	25.0; 28.5	27.2; 26.9
1	25.8; 20.5	22.6; 22.6	21.7; 28.1	24.0; 28.5	26.2; 26.9
2	28.1; 22.7	28.1; 21.7	23.4; 23.4	21.0; 28.5	23.2; 26.9
3	28.5; 25.0	28.5; 24.0	28.5; 21.0	22.3; 22.3	18.2; 26.9
4	26.9; 27.2	26.9; 26.2	26.9; 23.2	26.9; 18.2	19.1; 19.1

It is routine to verify that those presented in the text constitute the Nash equilibria of these games.

C Additional Tables

Table C1. Investment levels, winning bids and costs. Auction-level data.

	i_I	i_C	b_{it}^W	c_{it}^W
B	0.494**	0.567**	-0.252	3.428*
	(0.215)	(0.239)	(3.004)	(1.903)
c_{it}			0.620***	
			(0.030)	
$B \times c_{it}$			-0.030	
			(0.042)	
Constant	2.492***	0.981***	29.98***	35.46***
	(0.152)	(0.169)	(2.111)	(1.346)
Observations	1080	1080	1080	1080

Notes. This table reports two-ways linear random-effects models that account for dependency within rematching groups and within individuals (standard errors in parentheses). The dependent variables in Columns 1-4 are, respectively, i_I (the investment level of the Innovator in the innovation stage), i_C (the investment level of the Challenger in the innovation stage), b_{it}^W (the winning bid in the auction stage), and c_{it}^W (the production cost of the winning bidder in the auction stage). B is a treatment dummy, c_{it} is the production cost assigned to the bidder in the period. Significance levels are denoted as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

Table C2. Investment in the innovation stage of *PING*, data at the subject-level.

Panel A: Descriptive statistics

	B	*S*
i_{it}	2.280	1.743
	(1.455)	(1.347)
$k(i_{it})$	7.311	4.848
	(5.910)	(5.039)
Observations	2160	2160

Panel B: Parametric analysis

	All
B	0.537^{**}
	(0.220)
Constant	1.743^{***}
	(0.156)
Observations	2160

Notes. Panel A reports descriptive statistics (mean and standard deviation in parentheses). Panel B reports a two-ways linear random-effects model that accounts for dependency within rematching groups and within individuals (standard errors in parentheses). i_{it} and $k(i_{it})$ are, respectively, the innovation level and the associated cost borne by subject i in period t in the innovation stage of *PING*. The dependent variable in Panel B is i_{it}. *B* is a treatment dummy. Significance levels are denoted as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.