Mechanism Design for Designs: Principals with Taste

Matthew Mitchell*

February 9, 2016

Abstract

This paper considers a moral hazard problem where the output is a design. The design refers to something whose value is subjective, possibly difficult to describe, and can be returned if the principal views it as unsatisfactory. Correlation of the principal's signal with the his true value plays a role, in contrast to standard principal agent problem; experience goods are different from credence goods. The principal's ability to forecast value corresponds to a notion of taste for the principal that distinguishes taste from judgment. The importance of "taste" impacts the cost of contracting as well as the type and number of agents contracted with. For examples where taste is relevant, uncertainty in the outcome may make the incentive contract less costly. Negative correlation between principal and agent's signals can sometimes be valuable; this can be interpreted as a value in "bad" taste of the principal. The results can be applied to the use of termination in dynamic relational contracts.

*University of Toronto. Thanks to Heski Bar-Isaac, Jim Campbell, Rahul Deb, Matt Grennan, Marcin Peski, Nicolas Sahuguet, Florian Schuett, and participants at the Toronto Theory and IO Bag Lunches as well as the 2015 CIRANO-CIRPÉE Conference on Industrial Organization at HEC Montreal and the Mannheim Workshop in Innovation Economics for comments.

1 Introduction

Technological change that comes through new products is driven by the process of design. Design is, more broadly, a quantitatively important part of the economy. Narrow measures of the contribution of "creative industries" to world output exceed four percent and are fast growing; authors have argued that a variety of "knowledge" work outside of the creative industries might be viewed as part of the design process, leading to estimates that employment in creative intensive jobs greater than 30% (Florida [2004]), a doubling in 50 years. Growing rates of patenting indicate an increase in the rate of change in the production of new designs.

While important, design work may be difficult to monitor directly because effort is mental. Evaluating designs is often coarse and subjective, without an obvious way to pay based on a clear realization. When a new building is built or advertising campaign run, it can be difficult to assess, even ex post, the contribution of the project to output.[1] Potential forms of output are difficult to describe, since the details of what would be contained in a good design are unknown ex ante; this is what necessitates the need for a designer.[2] Moreover the designer may be better able to understand a good design than the principal.

Motivated by the design process, this paper studies a principal-agent model where the principal has a *non-verifiable* signal of the action of the agent or agents. The contract can specify not only payments, but, unlike many labor contracts where the work is complete, an allocation of the design: if the principal reports that he does not like a particular design, he may not keep it. This contract term is used in some cases in practice with design

[1] When the military asks for designs for a new piece of military equipment, it is difficult to assess the value added of that equipment. One imagines that the design of the exterior of the iPhone is relevant to its value, but it is hard to determine the benefit of that vis a vis other unique characteristics of the product. When a marketing firm is hired to design an ad campaign, it is difficult to determine even ex post the financial impact of the ad campaign. Even road construction projects often include a design component that may be rated, subjectively, and the ratings used to decide which firm undertakes the project (Takahashi [2014]).

[2] In fact, even in cases like the well documented case of the prize for the naval chronometer, a detailed description of the design's requirements was insufficient. Upon seeing a design that seemed to meet the criterion, the principal (the "Board of Longitude") denied that the design was sufficient; the principal did not believe the previously unknown innovator nor the simple design could possibly solve the very hard problem at hand.

contracts and is often called *right of refusal*. Right of refusal is useful in this context, and is related to the principal's ability to assess his own value, an element not relevant in canonical agency models. The main goal is to study a model where taste matters: the relationship between the principal's signal, their value, and the agent's value are interrelated in a way that can be interpreted as strength and nature of taste and judgment of the principal.

In the standard model of a principal and agent where the principal's signal is verifiable, the relevant informativeness of a signal is its correlation with the agent's action. In that textbook model it is irrelevant whether the payoff for the principal is exactly the signal, (for instance in the case where the signal is realized output of the agent's effort or the realized utility of an experience good), or the payoff is exactly the agent's action, so that the signal has no informativeness given the incentive compatible action (a credence good). Practically, the principal is merely trying to assess the magnitude of inputs; measuring output is just one possible route to that goal. Given the incentive compatible action, output is irrelevant for incentives, and should be assigned to maximize total surplus, without any regard to incentives. As a result, right of refusal, and more generally assessing *output conditional on inputs*, has no place in a standard model of incentive contracting.

When the principal's signal is private, however, the allocation of the output can be useful in generating incentives for the principal to report his signal truthfully. In particular, it is useful to the extent that the principal's signal has information content about the value of the output, so that the principal will be reluctant to refuse a project that he assesses to be valuable. This motivates the use of right of refusal, and allows the principal's ability to assess output to play a role.

The first results fix the relationship between the signal and the value of the design, and discuss the value of the precision of the signal. In many cases there is value in coarser signals when right of refusal is used. This is a broad sense not surprising, since the model assumes the signal is private information of the principal, and more *private* information is often not socially desirable. Coarser signals are not, however, beneficial in with private signals but *without* right of refusal as in MacLeod [2003]; there a coarser signal always makes the principal worse off. The value of coarseness is specifically on the high end of signals: when the project is kept only in some (good) states, the principal prefers, among states where the project is to be retained with probability one, to be unable to further distinguish between the states. In other words, conditional on liking the project sufficiently well to keep it for sure, the principal does not want precise information on exactly how much he "loves" the project.

The majority of the results are comparative statics on the impact of changing the degree of the principal's ability to assess the value of output. The model has a parameter that allows the signal to vary from a credence good to a pure experience good. The credence good case, where the principal cannot judge output at all conditional on the agent's input, is associated with the principal's noisy signal reflecting his lack of *judgment* about the design. In the experience good case, the signal is fully reflective of value ex post, and the noisy signal reflects the idiosyncrasy of the principal's *taste*. The closer the signal is to an experience good, the greater is the use of right of refusal is useful. This fits the casual evidence that it is common for commissioned artwork to include a "right of refusal" on the buyer side: if the buyer does not like the final product, he does not pay the full amount and, importantly, does not keep the artwork.[3] On the other hand, right of refusal is less common in a variety of design contracts that involve subjective evaluation for the principal, but greater expertise of the principal, for instance in architecture contracts.[4] When the agent is the expert on the value of the object, so that the signal provides relatively little information about the true value of the design, right of refusal is not as useful a contract feature.

Comparative statics on experience versus credence goods can be interpreted as studying strength of taste relative to judgment. Another comparative static alters the agent's value as a function of the principal's signal. When the values go in the same direction, payoffs are positively correlated; when values go in opposite directions, they are negatively correlated. In some cases the principal benefits from negative correlation. To the extent that the agent is the expert and the "connoisseur" would have similar tastes to the artist, this suggests that sometimes "bad taste," meaning the opposite of the artist or a connoisseur, is good for contracting costs. This conforms

[3] http://www.finearttips.com/2009/11/commissioned-art-%E2%80%93-tips-to-make-it-a-success/, http://thepracticalartworld.com/2011/11/21/how-to-build-a-contract-for-commissioned-artworks/ . A popular story of a commission being "refused" is the destruction of Diego Rivera's fresco "Man at the Crossroads." While clearly showing the possibility of refusal for art, it does not strictly fit the model in the sense that Rivera was paid the full commission even though it was destroyed. On the other hand, destroying the fresco clearly lowered Rivera's utility, which would fit the notion that the principal takes an action which lowers both his payoff and the agent's payoff in some states.

[4] http://www.raic.org/practice/contract_documents/index_e.htm

to experience of frictions between patrons and artists where artists think the work is good but the patron does not agree. In the model this value of bad taste can only happen when the principals tastes are sufficiently weak.

The comparative statics on the principal's ability to assess output also provide insight into the use of multiple agents. Tournaments are a well known solution to verifiability on the principal's side: the principal hires multiple agents and the commits to pay some agent, and therefore has no incentive to misreport the winner, as for instance in Lazear and Rosen [1979] and Malcomson [1984]. In the usual tournament model the principal benefits directly from both agents effort, for instance where both are selling a product or producing output, which the principal keeps. Designs are inherently somewhat indivisible.[5] In these cases there is the matter of duplication: the designs that are not chosen in the contest are, essentially, wasted; they have value in generating incentives, but the effort is not directly beneficial to the principal. This tension between, on the one hand, competition as an incentive device and, on the other hand, redundant effort in competition, is a fundamental feature of the design process with multiple agents.

The comparative statics show that using two agents is more likely for cases where the principal cannot assess the value of output, as in a credence good, but less valuable for experience goods. In other words, one might commission fine art from a single agent, when the goal is simply to enjoy the art, or an exclusive supplier in a vertical relationship, where expertise is strong on the buyer's side, but run a tournament for situations like architects and marketing firms, where the principal imagines that the designer is the best assessor of the quality of the design.[6] This force toward the use of multiple agents is not because of the usual channel of correlated signals in models of "second opinions."

An alternative interpretation of the single agent case is as one stage of a fixed-review contract as in Levin [2003] and Fuchs [2007] with both moral hazard and adverse selection. The ability to value output is analogous to learning about the agent's permanent type from the current review period signal. Adverse selection might therefore be helpful in solving moral hazard: if a good outcome signals a good draw on the agent's type, it is costly for

[5] While one can, to some extent, take the best features from multiple designs, it is certainly costly to aggregate the work of multiple designs.

[6] Contests between agents are common in both marketing and architecture, where several designs are solicited. More generally, product innovations are pitched to firms before they are produced and sold.

the principal to misreport the signal and fire the agent.

Literature

The fact that key features of design make it unusual relative to standard examples has been identified by Caves Caves [2003]. In his words, “Consumers’ reaction to a product are neither known beforehand, nor easily understood afterward.” Since the model assumes the principal’s signal is private, a natural starting point is MacLeod [2003]’s model of a principal-agent problem with a privately observed signal. The model here differs by allowing the contract to specify whether the agent’s output will be allocated to the principal, which allows a role for the principal’s taste to play a role, and by allowing the principal to contract with multiple agents. In MacLeod [2003], as well as many other papers studying subjective evaluation, subjectivity is identical to unverifiable or private. This paper describes subjectivity in more detail, taking private information as one ingredient in subjectivity. In particular, the model draws a distinction between private signals of the principal that reflect the principal’s preference conditional on the agent’s action on the one hand, and private signals that reflect merely the principal’s inability to identify the action of the agent on the other. I argue below that the former might be a more natural definition of subjectivity.

The model in this paper is also related the model of up-or-out promotion in Kahn and Huberman [1988]. In that model, unlike MacLeod [2003], the principal can refuse to continue to employ the agent if productivity is low. The costly action of the agent is often taken to be some investment that has a persistent impact on output, so that the cost to the employer of a breakup is higher when the investment is made. In this paper the key issue is the slope of E in s; the principal’s ability to assess output, which is fixed in the typical up or out promotions setup. Moreover, In Kahn and Huberman [1988], the principal is assumed to perfectly assess the value of the worker, unlike here.[7] The model in this paper highlights the role of the principal’s ability to assess the value of output, conditional on inputs, in features of the optimal contract such as the use of right of refusal and multiple agents.

Right of refusal is similar to a money back guarantee. Money back guarantees have been extensively analyzed in double-moral hazard environments

[7] Many other papers consider up-or-out structures including Ghosh and Waldman [2010], who consider the informational content of promotions, and Zabojnik and Bernhardt [2001].

(see Mann and Wissink [1988]) or as signals of quality (Lutz [1989]). Here the role is different: the principal provides no effort to the project, and there is nothing to signal. Similarly to the double moral hazard case, however, there is private information on both sides. Moreover this paper describes how the contract varies with the principal's ability to assess output.

In that sense the multiple agent setup is very much in the spirit of tournament models, such as Lazear and Rosen [1979] and Malcomson [1984]. Those models incorporate multiple agents with an information structure that is the polar opposite of moral hazard in teams (as in Holmstrom [1982]), where the principal sees a signal of team output but no signal of how that effort is distributed across team members. As in the tournament literature, the principal here gets a signal of relative inputs, but no signal of the absolute level of team effort. As a result, the structure with multiple agents is a sort of tournament.[8]

Paper Organization

Section 2 introduces the general model with one agent. It builds directly on MacLeod [2003], where a principal contracts with a single agent, but with the added possibility of "right of refusal." Then further results and comparative statics are developed in the two action, two outcome case, in section 3. The two agent case is studied in section 4. Finally, the model's implications for relational contracting (in 5) and innovation prizes (in 6) is discussed.

2 The Model

Throughout, an agent takes an action denoted by $a \in R$, which is unobserved by the principal, at convex cost $c(a)$. The principal observes a signal s of the agent's action, which is unverifiable, because it is subjective. Take the signal to be distributed over a finite set; this follows MacLeod, and conforms to the notion that subjective views of a design project are unlikely to be very fine-grained, but is not essential for the results. The signal is distributed according to $f(s|a)$.

[8]Our model with two agents is also related to the problem in Che and Gale [2003]. Their quality is unverifiable, but, in strong contrast to this paper, is certain and observed by both parties, in the spirit of Hart and Moore [1988], and therefore not subjective.

2.1 No reallocation of output: MacLeod (2003)'s Money Burning Contract

MacLeod allows for the principal to commit to destroy (burn) resources: the agent is paid $w(s)$, and an amount $b(s)$ is burned. The principal is risk neutral and allocated the value of the action a regardless of the state, so it does not impact his reporting incentives. The agent values the wage by the concave $u(w)$ and has outside option normalized to zero. With inefficient money burning, ex post renegotiation is beneficial; renegotiation is discussed at the end of section 3. In some cases, right of refusal is beneficial in solving renegotiation issues.[9]

The cost to the principal of implementing an action a is

$$\begin{aligned} C(a) &= min_{w(s),b(s)} \sum (w(s)+b(s))f(s|a) \qquad (1)\\ & s.t. \\ a &\in argmax_{\hat{a}} \sum u(w(s))f(s|\hat{a}) - c(\hat{a}) \\ 0 &\leq \sum u(w(s))f(s|a) - c(a) \\ s &\in argmin_{\hat{s}} w(\hat{s}) + b(\hat{s}) \end{aligned}$$

The solution must have $w(\hat{s}) + b(\hat{s}) = max_s w(\hat{s}) + b(\hat{s})$, and therefore the cost is equal to the maximum resources expended. This leads, under the usual monotonicity condition on F, to the planner using just two levels of payment.[10]

Proposition 1. *(MacLeod 2003). Suppose that F has the monotone likelihood ratio property, $\frac{df}{da}/f$ increasing. Then the optimal contract has $b = 0$ (and therefore a constant wage) for all but the lowest value of s.*

Section 3 studies a model with only two signals; this both fits the coarse nature of tastes for design and, in many cases such as this one, may be without loss, in the sense that the optimal policy uses at most two outcomes for the contract. It is immediate from the fact that the principal is indifferent to all reports that a coarser signal must make the cost weakly higher; whatever is optimal given the coarser signal is incentive feasible for a finer

[9]Dynamic models like Levin [2003] and Fuchs [2007] allow the burned resources to be interpreted as lost future contracting opportunities when the relationship breaks up. Section 5 draws out this relationship in more detail.

[10]Proofs are in the Appendix when not included in the text.

signal, since the payments are the same for all signals in either case. It turns out that coarseness of the signal may be valuable when the principal chooses to discard a design; with only money burning, however, a coarser signal, in particular at the bottom of the distribution, makes the principal worse off.[11]

2.2 The Role of Taste and Judgment with Right of Refusal

One can interpret s as a measure of output; however, in both the standard principal agent model and the model of MacLeod, s is fundamentally an assessment of the agent's *input*. The role of assessing *output* becomes clear when the principal can choose to refuse the design as a way to satisfy the constraint on reporting s. This feature matches the common "right of refusal" included in some design contracts. One goal of the paper is to understand when such a feature might be useful to include in the contract.[12] One can interpret it either as the design being destroyed, or that the design is reallocated to the agent, who has a low value for discarded designs.

The contract calls for the principal to discard the project with probability $p(s)$.[13] In order to understand the impact of this choice on the principal's payoff requires describing the planner's valuation of designs conditional on s, which is the sense in which taste matters. Let $E(s,a)$ denote the principal's expected value of the design, given recommended action a and signal s. It is assumed that $E(s,a)$ is increasing in both arguments. On the other hand, the principal is assumed to get zero value if the design is discarded.

In constructing the cost function, the costs of burning are taken relative to the case where the planner keeps the project, so that his gross expected value, before costs, is $\sum E(s,a)f(s|a)$. This expected payoff is assumed to be linear in a and so can be normalized as $\sum E(s,a)f(s|a) = a$. Costs of discarding the project are included in the cost function $C(a)$, so the principal's problem can be reduced to choosing a recommended action a to maximize

$$\sum E(s,a)f(s|a) - C(a)$$

[11] For cases that do not involve the bottom state the coarseness has no effect since the principal takes the same action in those states.

[12] This seems at least as reasonable as committing to money burning in some examples, where the output is a well defined object which can be returned.

[13] The contract without randomization is discussed below.

where $C(a)$ is the minimized cost of delivering action a; the objective is simply $a - C(a)$. The cost of implementing a is

$$\begin{aligned} C(a) &= min_{w(s),b(s),p(s)} \sum \left(p(s)E(s,a) + w(s) + b(s)\right) f(s|a) \qquad (2) \\ & s.t. \\ a &\in argmax_{\hat{a}} \sum u(w(s))f(s|\hat{a}) - \hat{a} \\ 0 &\leq \sum u(w(s))f(s|a) - c(a) \\ s &\in argmin_{\hat{s}} p(\hat{s})E(s,a) + w(\hat{s}) + b(\hat{s}) \end{aligned}$$

Note that, since $E(s,a)$ could be negative for some s, if the project has negative value contingent on the signal, so that the term $p(s)E(s,a)$ could be negative.

If the design can be refused, the details of the mapping E are relevant because they impact the incentive constraint describing $C(a)$. The interpretation of E is related to whether the principal or the agent has better information about the principal's valuation of the design. For example, if one interprets the design as a credence good, where the value is known by the agent but the principal has no information, s is not informative and therefore $E(s,a) = a$. In that case it is immediate that the problem (2) is equivalent the the one in (1), i.e. similar to the one in MacLeod (2003); there is no value in a right of refusal when the principal knows that the agent is the expert, and the signal conveys no information given the recommended action. The case where E is constant in s is discussed explicitly below, but in the rest of this section $E(s,a)$ is assumed to be strictly increasing in s, to differentiate from MacLeod [2003]. For instance, a natural case is $E(s,a) = s$, so the signal is exactly the valuation of the principal of the object, which is consistent with the output being a pure experience good. This is often the interpretation in the canonical model principal-agent model, where the agent produces output for the principal stochastically by exerting effort. [14]

[14] E and f appear in this problem and therefore are treated as primitive in this section. However one could think of modeling value explicitly. Specifically, suppose that value v for the principal, which is not contractible, depends on the effort of the agent according to $g(v|a)$. Signals are associated with some combination of value (to the extent that the principal seeks to assess the final product) and inputs (to the extent that assessment is about inputs) by $h(s|a,v)$. Therefore

$$f(s|a) = \int h(s|a,v)g(v|a)dv$$

The planner can use an decreasing $p(s)$ (so the probability of keeping the design, $1 - p(s)$, increases is s) to help solve the reporting incentive: the planner may pay less for lower s, but on the other hand gets to keep the design less often. This trade off may be attractive if the true state is low, but not if the true state is high. The reverse is never incentive compatible:

Lemma 2. *$p(s)$ is weakly decreasing*

Discarding the project in states with $E(s,a) < 0$ is costless, and in fact is the optimal choice without reporting constraints. Therefore if $E(s,a) < 0$, $p(s) = 1$. The only reason for refusing a positive value, ex post, is because the planner is using the refusal to get incentives for a higher state. Lower states are not a problem: there is always a rate at which the planner could pay (through burning money) to keep the project more often, improving the total payoff for s and such states less than s would not find the trade-off attractive. As a result, each project worth keeping ($E(s,a) > 0$) must either be kept for sure, or be limited by a higher state s' that is indifferent to reporting s. This idea is formalized in the following lemma:

Lemma 3. *Each s with $E(s,a) > 0$ has either $p(s) = 0$ or*

$$p(s)E(s',a) + w(s) + b(s) = p(s')E(s',a) + w(s') + b(s')$$

for the smallest $s' > s$ with $p(s') < p(s)$. Further, for no $s' < s$ does the equality hold.

Since every contract is strictly preferred to misreporting and getting a strictly lower $p(s)$, one can always, if $p(s) < 1$ and $b(s) > 0$, raise $p(s)$ and lower $b(s)$, keeping the payoff the same for truthful reporting of s and maintain incentive compatibility for lower types (who strictly prefer not reporting s). This makes the contract for s strictly worse for $s' > s$. In other words resources are only burned if the planner is certain to discard the project:

Corollary 4. *If $0 < p(s) < 1$ then $b(s) = 0$*

---

and

$$E(s,a) = \int vh(s|a,v)g(v|a)dv$$

An explicit discussion of this interpretation is developed in the next section.

The corollary implies that $p(s) < 1$ for at least some s if $a > 0$, since otherwise $b(s)$ is zero for all s and the only incentive compatible contract for the planner has $w(s)$ constant, and therefore cannot be incentive compatible for the agent.

Whenever there are some states where the project is burned inefficiently (i.e. $E(s,a) > 0$ and $p(s) > 0$), there is the potential for improvement from coarsening states. The next result shows that it must be the case that the principal would (weakly) benefit from a coarsening that combines states where $p(s) = 1$. Let $P = \{s|p(s) = 0\}$ for the problem in (2). Denote the conditional probabilities on P under f by f^c. The optimal contract is constant on P; if one makes information coarser, so that the principal cannot distinguish between P, the policy remains feasible; moreover, it is incentive compatible since for the minimum s in P:

$$\begin{aligned} p(s-1)E(s,a) + w(s-1) + b(s-1) &\leq p(s)E(s,a) + w(s) + b(s) \\ &< p(s)\sum_{x\in P} E(x,a)f^c(x|a) + w(s) + b(s) \end{aligned}$$

Therefore it is immediate that:

Proposition 5. *Let $\tilde{C}(a)$ be the cost under the coarsening of information such that the principal cannot distinguish states in P. Then $\tilde{C}(a) \leq C(a)$.*

To understand how coarseness could be useful, and to see how P might contain more than one state, consider the following example where u is linear. Under linearity, absent the principal's private information on s, the outcome is first best. Let $s \in \{0,1,2\}$, with, for recommended action a, $E(0,a) < 0 < E(1,a) < E(2,a)$. Let $w(1)$ and $w(0)$ be defined by

$$\begin{aligned} a &= argmax_a f(0|a)w(0) + (f(1|a) + f(2|a))w(1) \\ 0 &= \sum w(s)f(s|a) - c(a) \end{aligned}$$

and suppose that, for such $w(1)$

$$E(1,a)f^c(1|a) + E(2,a)f^c(2|a) + w(0) = w(1)$$

As a result, the first best is attained by using $w(0)$ in state 0 and $w(1)$ in states 1 and 2. However if states 1 and 2 can be distinguished the contract no longer satisfies the reporting constraint for state 1, since

$$E(1) + w(0) < E(1,a)f^c(1|a) + E(2,a)f^c(2|a) + w(0) = w(1)$$

As a result it is impossible for the contract to continue to have $p(s) = 0$ for $s = 1, 2$ even with $p(0) = 1$. The first best cannot be implemented.

Being able to discern s may therefore be counterproductive. The principal can do better if it can commit not to become "too well" informed; it does nonetheless benefit from information conditional on the contract. This is a sense in which a lack of taste, in the sense of being able to discern states, can have benefit. The ability to judge output conditional on s, however, is always useful for the principal, in the sense that a stronger relationship between s and the final payoff lowers costs:

Proposition 6. *For any s such that $0 < p(s) < 1$ there exists $\kappa_h, \kappa_l > 0$ with $\sum_{s>\bar{s}} \kappa_h f(s|a) = \sum_{s<\bar{s}} \kappa_l f(s|a)$ such that, if $E(s,a)$ is replaced with $\hat{E}(s,a)$ such that $\hat{E}(s,a) = E(s,a) + \kappa_h$ for all $s > \bar{s}$ and $\hat{E}(s,a) = E(s,a) - \kappa_l$ for all $s < \bar{s}$, the principal's cost is lower.*

In the standard model, only the correlation between signal and actions, embodied in $f(s|a)$, matters for incentives; here, taste or judgment about output matters. In particular, stronger taste or better judgment, as reflected in a more informative signal that is more strongly related to value, can lower cost for the principal. It may be especially important in the sense that more information, in terms of distinguishing states, may be detrimental.

3 A 2x2 Model

In this section a further connection between f and E is developed through a common shock in the context of the "2×2" structure common in moral hazard problems: there are two actions, normalized to be some fixed $\bar{a}$ and 0, and two signals, 0 and 1.

Actions are linked to value v for the principal according to $v = a - \lambda\epsilon$, where $0 \le \lambda \le 1$. The signal s is a noisy and coarse signal of value, which takes on the value 1 if $v - (1-\lambda)\epsilon > 0$, and 0 otherwise. In other words the principal observes the signal one if and only if

$$a \ge \epsilon$$

where ϵ is distributed according to the symmetric, mean zero distribution $F(\epsilon)$; therefore the signal one is observed with probability $F(a)$, denoted by shorthand F in this section, if the agent chooses a, and probability 1/2 otherwise. The relationship between a and the signal, therefore, is governed by F.

The case of binary signals is especially interesting in the context of subjective evaluation. First, it is natural to imagine that, especially in the case where the agent is the expert, but more generally with design projects where assessment is subjective, that it may be possible for the principal to decide whether or not the planner likes the design more than the outside option, but much harder to quantify the differences between designs. Second, from MacLeod [2003] and the results above, subjective evaluation leads to coarse (often binary) policies even for richer signal spaces. Finally, binary signals generates a tournament-like signal structure that allows for a direct comparison to the tournament solution to non-verifiability in Malcomson [1984], which is done in the next section.

One natural interpretation is that $a - \epsilon$ measures the principal's assessment of the agent's output, relative to an outside option of zero, so that the project appears better than the outside option if $a > \epsilon$.[15] A signal of 0 may or may not be strongly indicative of the value of the project being less than the outside option, conditional on a; $E(s, a)$ is given by

$$\begin{aligned} E(1, a) &= a - \lambda \frac{\int_{x<a} x f(x) dx}{F} \\ E(0, a) &= a - \lambda \frac{\int_{x>a} x f(x) dx}{1 - F} \end{aligned}$$

For shorthand, let $E(s, \bar{a}) = E(s)$, since the focus will always be on eliciting high effort from the agent. The slope of E in the signal, conditional on $\bar{a}$, is governed by λ, while λ has no impact on the relationship between signals and actions; the fraction λ of the shock is relevant to output. When $\lambda = 0$, the planner truly values projects based on the effort, but evaluates them with a noisy indicator of that true quality. The agent is an expert, in the sense that the subjective evaluation is only erroneous noise. Variance of ϵ denotes a lack of *judgment* on the principal's part: he cannot discern the good from the bad, as the final output is a credence good. In many design projects, the principal must choose a design, but knows that their option is less informative than that of the agent which they hired. When $\lambda = 1$ the noise reflects the actual underlying idiosyncratic taste of the planner, as might be natural in an art

[15] This could be due either to randomness in evaluation of the project or the outside option.

project: the noise reflects noise in the true value of the project. In that case variance in ϵ denotes not a lack of judgment, but rather a strong subjective *taste*. This assumption is the classical one in principal-agent models where the agent gives effort that leads to stochastic output for the principal, but in the typical model with contractible signals, all that matters is the correlation between a and s (which is determined solely by the distribution of ϵ); true taste (and therefore λ) plays no role.

Define $\bar{\lambda}$ such that, given the signal is unfavorable to the project with effort $\bar{a}$, the expected net value is zero, i.e. $E(0) = 0$. This value is always unique and interior since, for $\lambda = 0$, $E(0) = a>0$, and for $\lambda = 1$, $E(0) < 0$, and $E(0)$ is strictly decreasing in λ.

3.1 No Reporting Constraints

First, as a benchmark, consider the case with no reporting constraints. In that case the optimal payments for an agent that works a are either the w and l that satisfy the IR and IC constraints. The IC constraint for a reduces to

$$u(w_1) - u(w_0) = 2c/(2F - 1)$$

and participation

$$Fu(w_1) + (1 - F)u(w_0) = c$$

For the solution to this problem, let $\omega = w_1 - w_0$.

Without reporting constraints for the principal, if $\lambda < \bar{\lambda}$, the planner always takes the project, since it is better than not, and the payoff is

$$\bar{a} - Fw_1 - (1 - F)w_0$$

If, on the other hand, $\lambda > \bar{\lambda}$, the planner only takes the project if the signal indicates it has value, and the payoff is

$$\bar{a} - (1 - F)(\bar{a} - \lambda E(\epsilon|\epsilon > \bar{a})) - Fw_1 - (1 - F)w_0$$

Define $C^{VS}(\bar{a})$ so that this payoff is $\bar{a} - C^{VS}(\bar{a})$.

3.2 Cost With Reporting Constraints

The cost of implementing effort a is

$$\begin{aligned} C(\bar{a}) = &\quad min_{p_i,w_i,l_i,b_i} \begin{pmatrix} F(w_1 + p_1E(1) + b_1)+ \\ (1-F)(w_0 + p_0E(0) + b_0) \end{pmatrix} && (3) \\ & \qquad\qquad s.t. \\ Fu(w_1) + (1-F)u(w_0) - c \geq &\quad .5u(w_1) + .5u(w_0) \\ 0 \leq &\quad Fu(w_1) + (1-F)u(w_0) - c \\ w_1 + p_1E(1) + b_1 \leq &\quad w_0 + p_0E(1) + b_0 && (4) \\ w_1 + p_1E(0) + b_1 \geq &\quad w_0 + p_0E(0) + b_0 && (5) \end{aligned}$$

Note that $b_1$is unnecessary, since the planner can always use w_1 at the same cost, and only improve incentives. If $p_1 = p_0$, then one of the reporting constraints is sufficient for the other. Since p is weakly increasing, the only other possibility is that $p_1 > p_0$. It is immediate that in that case, only one of the reporting constraints can bind, and it is the reporting of $s = 1$. Therefore the last constraint does not bind.

3.2.1 Principal with Strong Idiosyncratic Tastes

Suppose that $\lambda > \bar{\lambda}$. The principal will always keep a design that has signal $s = 1$, and refuse one with $s = 0$; the refusal helps generate incentives since it is particularly costly to refuse a design that appeals to the principal. Therefore the planner's problem is characterized as:

Proposition 7. *Suppose* $\lambda > \bar{\lambda}$. *Then* $p_1 = 0$ and $p_0 = 1$, and $C(\bar{a}) = C^{VS}(\bar{a}) + max\{(1-F)(\omega - E(1), 0\}$

When $\lambda > \bar{\lambda}$ the contract uses right of refusal to achieve the incentives: if the planner wants to claim the signal is bad, he cannot enjoy the output. Therefore higher λ, which means that the signal is more closely associated with the principal's payoff, improves outcomes.

Large λ is consistent with the principal being the "expert" in terms of assessing the principal's payoff. One reason this would happen would be because the principal's payoff is, in fact, his subjective affinity for the design. In other words, for truly subjective cases, in other words matters of subjective taste, λ is high, right of refusal is a useful part of the contract.

To see how the planner might do better with coarser information, as described generally above, consider the case where $E(1) > \omega$, so that $C(a) = C^{VS}(a)$ Now suppose that the state "1" is divided into two states, 1 and 2, where $E(2) > \omega > E(1) > 0$. Doing as well as in the coarse case requires that the design always be kept in the two states, since it has net value; however, doing so implies that the extra payment must be at least ω in both states, which is not incentive compatible (without burning resources) in the state 1. Therefore the planner, in the finer example, will need to employ some combination of burning resources or the project to achieve incentive compatibility, meaning costs are greater than C^{VS}. Again coarseness might be viewed as arbitrariness on the principal's side, and might be beneficial.

3.2.2 Principal with Little Taste

Now suppose that $0 < \lambda < \bar{\lambda}$,so the principal is sufficiently uninformed that the high effort project is better even when the signal is not favorable to it. The cost function is:

Proposition 8. *Suppose* $\lambda < \bar{\lambda}$. *Then*

$$C(a,0) = \begin{cases} C^{VS}(\bar{a},0) - (1-F)\omega\frac{E(0)}{E(1)} & \text{if } E(1) \geq \omega \\ C^{VS}(\bar{a},0) + (1-F)(\omega - E(1) + E(0)) & \text{if } E(1) < \omega \end{cases}$$

In this case refusal is typically probabilistic; since the principal knows that, given a, the output has positive value, the contract only refuses to the extent necessary to get incentives, trading off the lost value against the incentive benefit. If the principal had to choose a refusal level non-stochastically, sometimes it would be worth refusing the project, and sometimes not. The lower is λ the less is the value of refusal as an incentive device, and the greater is the value of the project in the good state; therefore, right of refusal would be used only for sufficiently high λ, and no refusal would be used when λ was low enough. Therefore the model in that case specifically delivers the notion that sometimes right of refusal is used (specifically for high λ) and sometimes not.

3.3 Taste and Cost

In the analysis that follows the key object of interest is the extra cost incurred above and beyond the case with verifiable signals, denoted by $\Gamma = C(\bar{a}) - C^{VS}(\bar{a})$.

From the previous section it is immediate that

Corollary 9. $\partial\Gamma/\partial\lambda < 0$.

Strong taste, in the sense of higher λ, lowers the informational cost of contracting.

In general, fixing $F(a)$ but increasing $E(1)$ also makes incentives easier to achieve, for instance by a mean preserving spread on the set $(-\infty, -\bar{a}] \cup [\bar{a}, \infty)$. Moreover the benefit of such a change in the distribution of ϵ is complementary with λ. In other words, principal's with strong taste, i.e. a high ability to judge output, benefit from having highly *variable* preference for output. This is consistent with the notion that patrons of the arts often have taste that is both strong and unusual (in the sense of not being very predictable).[16]

3.4 Substituting taste for ability

This section considers the relationship between λ and the value of the agent's quality. In some cases, having the ability to assess the value of output may substitute for the quality of the agent, in particular when λ is high enough.. In order to model different qualities of agents, consider raising $\bar{a}$, fixing $c(\bar{a})$. This can be interpreted as a higher quality agent in the sense the agent does more at the same cost to the agent. In the case of verifiable signals this is unambiguously positive. For principal's with high λ, the higher is λ, the less valuable is raising $\bar{a}$:

Proposition 10. *Suppose* $\lambda > \bar{\lambda}$. *Then* $\frac{\partial^2\Gamma}{\partial\bar{a}\partial\lambda} > 0$

Suppose that higher quality agents had a higher outside option $\bar{u}$. For $\lambda > \bar{\lambda}$ outside option is unrelated to Γ (i.e $\frac{\partial^2 C}{\partial\lambda\partial\bar{u}} = 0$) since Γ only depends on λ through $E(1)$ which is independent of the outside option. Therefore higher λ would be willing to pay less, in delivered utility $\bar{u}$, for quality than would a principal with lower λ in this range. In a sense quality of the agent is endogenously a substitute for lack of taste or judgment of the principal.

For $\lambda < \bar{\lambda}$ the same may not be true. While $E(1)$ is increasing in $\bar{a}$, so is $E(0)$; their difference is unambiguously *falling*. The ability to judge output (λ) is only unambiguously a substitute for agent quality when the ability is sufficiently strong, i.e. λ is high enough.

[16]The variability result can be shown in terms of variance for distributions like the Laplace distribution, which is analytically tractable.

3.5 Agreement and Disagreement Between Principal and Agent: When Bad Taste is Good

Up until now the agent's value of the object was set to zero. The level of this value can be thought of as a normalization, but importantly it implies that the agent's value of the object is be uncorrelated with the signal that the principal receives, even though the signal is positively correlated with the principal's value. This section considers both positive and negative correlation between the principal's signal and the agent's value of the object. The perhaps surprising result is that when $\Gamma > 0$, so that the friction studied here has some bite, negative correlation is valuable with right of refusal and $\bar{\lambda} > \lambda > 0$, but not otherwise. This clarifies a further role for taste beyond that in MacLeod [2003]. The interpretation is that frictions between the artist and patron, in the sense of a disagreement as to the value of the work, can actually help make contracting work better in the arts example, when taste is weak for the principal. Such disagreements are common even for artists of unquestionable quality.[17]

Correlation between principal and agent can also be seen as a central feature of taste. To the extent that the agent is an expert, negative correlation is a sort of bad taste: the principal likes what has low value. In that sense the model is one where "bad taste is good," but only in specific examples.

Suppose that the agent's expected value of the object is $A(s)$ when the signal is s and the action is 1. Consider how the cost C of providing incentives changes when $A(0)$ changes (marginally) away from the benchmark case of $A(0) = 0$.. Marginal changes to $A(1)$ are irrelevant for cost, since the object is strictly never offered to the agent after $s = 1$ and therefore won't be for small positive $A(1)$. Therefore one can think of changes to $A(0)$ as coming with corresponding changes to $A(1)$ that keep the agent's expected value of the object constant. Increases in $A(0)$ make value (perfectly) negatively correlated between agent and principal, and decreases make value positively correlated. In some cases this correlation impacts costs via Γ:

Proposition 11. *Suppose $\Gamma > 0$. If $0 < p(0) < 1$, then $\frac{dC(\bar{a})}{dA(0)} < 0$. Otherwise, $\frac{dC}{dA(0)} = 0$.*

To see the total impact of taste on cost, it is helpful to think about how C^{VS} changes with $A(0)$. For $\lambda < \bar{\lambda}$, the object is not allocated to the agent

[17] "To commission a painting from Rembrandt included a certain risk of dissatisfaction that was incongruous with the high prices he demanded." Crenshaw and van Rijn [2006]

in the low state with verifiable signals, and therefore $A(0)$ has no impact on C^{VS}; therefore the total impact of $A(0)$ on cost C is exactly the impact on Γ. Raising $A(0)$ lowers Γ and therefore lowers C.

For $\lambda > \bar{\lambda}$, and therefore $p(0) = 1$, C^{VS} falls with $A(0)$, as the allocated object is more valued by the agent who receives it, allowing wages to be lowered. However, since C is unchanged, the fall in C^{VS} is offset by an increase in Γ of exactly the same magnitude. Reporting constraints nullify the gains from trade benefit of $A(0) > 0$ with verifiable signals, if they bind.

One interpretation of "good" taste is that it is taste correlated with the expert designer. In that case, good taste corresponds to lowering $A(0)$. For moderate strength of taste ($0 < \lambda < \bar{\lambda}$), bad taste can be valuable: cost declines as $A(0)$ increases. A principal with strong tastes ($\lambda > \bar{\lambda}$) can also benefit from "bad" taste, but only to the extent that there are simple gains from trade when $A(0)$ increases, and the principal is doing as well as with verifiable taste. For the interesting case where $\Gamma > 0$, there is no benefit to bad taste for the principal with strong taste. The benefit is only when the right of refusal margin can be strengthened as taste gets worse. The model has a feature that conforms to experience of artists and patrons, as mentioned earlier for Rembrandt: some patrons with seemingly weak tastes none the less refuse projects even though the project is favored by the artist. Taste plays a role in those contracts differently from contracts with verifiable signals.

Inefficient allocations, such as money burning and refusing in the case where $\lambda < \bar{\lambda}$ for a single agent, are not renegotiation proof. There are two conclusions to potentially be drawn from this. For money burning, the natural solution is for the burned money to be delivered to a third agent who will be unwilling to renegotiate. This solution to renegotiation immediately leads one to wonder if a better solution than simply giving those resources to a third agent would be to have the third agent be truly an agent, potentially producing a design, making the structure more like a tournament. That case is considered explicitly in the next section.[18]

[18] Alternatively, one can imagine that a standard principal-agent model with only two agents in this case can only proceed if $\lambda > \bar{\lambda}$ and $E(1) > \omega$, so that the allocation is efficient ex post. Only agents with strong idiosyncratic taste can contract with an agent for a design; all others must bring a third agent into the contract, or cannot elicit high effort at all.

4 Tournaments with Multiple Agents

Right of refusal is related to a different form of discipline for a principal with a private signal: a tournament. In the classic tournament (Lazear and Rosen [1979], Malcomson [1984]), the planner can either report that one agent wins, or the other; in the symmetric case, the planner's payoff is identical regardless of the identity of the winner, and therefore reporting honestly is incentive compatible. This section allows for the possibility of hiring more than one designer. Contracts with competition for designs is common, for instance in architecture, marketing firms, or consulting firms. A relevant question is what the role is of tournaments versus pure right of refusal, for different underlying fundamentals.

Suppose the planner contracts with two agents, 0 and 1 to produce exactly one design. The designer picks a "winning" design and pays accordingly. The values of the two projects are

$$\begin{aligned} v_1 &= a_1 - \lambda\epsilon/2 \\ v_0 &= a_0 + \lambda\epsilon/2 \end{aligned}$$

The principal observes the signal one if and only if

$$a_1 - \epsilon/2 \geq a_0 + \epsilon/2$$

and the signal 1 otherwise. The structure nests the one agent model when $a_0 = 0$. Whereas before the model compared the outside option to the project, here the comparison is between designers, but without the outside option. Eliminating the outside option is a useful model because it avoids having multiple agents used because it simply generates more draws; this is certainly a reasonable motivation to use multiple agents, but not the one under study here.[19] If $a_0 = 0$ is interpreted as hiring one agent, the firm still gets the benefits of the same epsilon draws. Here the reason for contracting with multiple agents is purely due to incentive considerations. [20]

[19]The implications of multiple draws is clear, and not directly related to the information frictions studied here.

[20]In a sense, the model is the polar opposite of a moral hazard in teams problem. In that problem, both inputs are useful (and perhaps even essential); here only one input is useful ex post, since the principal uses one design. In team production, the principal gets a signal of combined inputs, but no measure of the division of inputs across team members, which generates a familiar free riding problem. Here the principal gets a signal that orders relative inputs, but no signal of combined inputs.

First suppose the action profile is (a, a). In that case the contract with verifiable signals is symmetric, and therefore $C(a, a) = C^{VS}(a, a)$. Moreover, since $C(a, 0)$ is exactly as in the two one agent case, the difference between costs with one and two agents can be written as

$$C(a, a) - C(a, 0) = C^{VS}(a, a) - C^{VS}(a, 0) - \Gamma$$

The incentive benefit of using a second agent, beyond what would be true with verifiable signals, is exactly Γ. By corollary 9, therefore, the incentive benefit of using a second agent declines with λ. In a sense, to the extent that right of refusal can be used, contracting with one agent acts like a tournament with the outside option; that tournament is only useful if it is credible that the outside option is a real competitor. By contrast, when actions are a for each agent, credibility is never a concern. Therefore principal's with an ability to judge output don't need to benchmark output to another agent's work; this is not because the second agent's work helps them to assess value, but rather because the second agent's work is a substitute, in terms of reporting s, for taste or judgment.

5 Dynamic Contracts

Levin [2003] and Fuchs [2007] study dynamic models with subjective evaluation. Levin focuses on essentially-static contracts that review the agent every period. Fuchs extends this arrangement to T period review contracts, and shows that no feedback is used during the review period, so the contract is again essentially static. In both cases, breakups can serve the role of money burning, to the extent that the relationship has value. In Fuchs, after T periods, the relationship is severed in the worst state (T consecutive failures), mirroring the worst state punishment in MacLeod [2003].

One way to interpret the static model here is as a fixed review period under a dynamic relational contract. The immediate conclusion is that principals who can value their employee's output, and potentially return it if it is unacceptable, will have lower turnover, but can have larger variation in wage payments across states. For instance, a patron can have a long term relationship with an artist, but not always reward the artist equally each period.

In many contexts, however, returning an employee's output is impossible; the work is done. In the rest of this section considers the possibility that the only recourse, besides possibly burning money, is to break up the relationship. Selection can potentially serve the function of taste: it can make the cost of a breakup higher when the signal is high.

To see this, consider the two state, two action model, but where agents are indexed by a productivity shifter θ that is unobserved to both parties, and where output is $v = a + \theta - \lambda\epsilon$. For simplicity, suppose the employment relationship lasts at most two periods; firing is possible in the first period, but not in the second, where the contracting environment is the same as in the earlier sections of this paper. Starting with the second period, the expected return for a given a, given a distribution Φ on θ for the employed workers, is denoted $V(a, \Phi)$. The following characterizes this value in terms of the expected F and E for extreme λ.[21]

Lemma 12. *Suppose $\lambda = 0$ or $\lambda > \bar{\lambda}$. Then $V(a, \Phi) = a + \int \theta d\Phi(\theta) - C_\Phi(a)$, where $C_\Phi(a) = C(a)$ with $F = F_\Phi = \int F(a+\theta)d\Phi(\theta)$ and $E(s) = E_\Phi(s) = \int E(s, a+\theta)d\Phi(\theta)$, ω_Φ solves the IR and IC constraints in expectation.*

Denote Φ_s the distribution after realizing state s in period 1.

Lemma 13. *Φ_1 first order stochastically dominates Φ_0*

Now focus on the first period problem as a static "one period review" as in Levin [2003] In addition to burning money, the second period value can be destroyed by severing the relationship, and therefore $V(a, \Phi_s)$ acts as the value lost from "right of refusal", i.e. $V(a, \Phi_s) = E(s, a)$. But

Proposition 14. *Suppose $\lambda = 0$ or $\lambda > \bar{\lambda}$. Then $V(a, \Phi_1) > V(a, \Phi_0)$.*

Now that V is characterized, the first period can be evaluated. The role of selection can be clearly seen for $\lambda = 0$, where right of refusal is worthless in the static model since $E(1) = E(0)$. In the dynamic model, however, the fact that $V(a, \Phi_1) > V(a, \Phi_0)$ implies that there breakup is in fact more efficient than money burning, since it is more costly to destroy a relationship after a good observation by the principal, and therefore the principal gets extra incentive to truthfully report as in the static model with $\lambda > 0$. Selection imparts taste into the principal.

[21] For $\lambda \in (0, \bar{\lambda})$ the characterization is also in terms of the expectation of $(1-F)E(0)$, however the proposition that follows cannot be verified for reasons similar to the ones in Proposition 10.

If one were to study a model with more than two periods, this effect would change over time. As the principal gets more information about the type, severing the relationship becomes more and more like money burning, as in the standard models. A complete discussion of dynamic contracting with moral hazard and adverse selection in this environment is beyond the scope of this paper, but the result here suggests that adverse selection can make the costs of a moral hazard problem lower, because of the incentives it generates on the reporting of states by the principal.

6 Subjective Innovation Prizes

One interpretation of the contract is as a prize for a design (or innovation) that can be awarded at the prize giver's discretion. The history of prizes, dating back at least as far as the prize for the naval chronometer, has many examples where assessment of the potential prize-winner was subjective, and incentives to renegotiate or renege on the prize award limit the effectiveness of the prize.[22]

Under the interpretation of the contract here as a prize, the prize granter needs to consider the foregone value when the prize is not granted. If the prize granter is "dispassionate" and has no preference over the outcome, then $E(s, a) = 0$. Many recent examples of prizes, such as the Gates Foundation vaccine prizes, are based on philanthropy; the prize granter wants the innovation to be used. In those cases, a lack of a prize could lead to a patent instead. The prize granter presumably cares about the deadweight loss from monopoly power, since that is a key reason for offering a prize in the first place. If deadweight loss scales with the market size for the innovation, then $E(s, a)$ is increasing in s: not giving a prize to a valuable innovation leads to a costly patent. As a result the prize organizer, even with a private estimate of the value of the innovation, may be willing to honestly award the prize, if it means avoiding the patent. Submissions for the prize which are deemed insufficiently high quality for the prize may still be patented. The reward is a hybrid between a prize and a patent, where the use of both is endogenous. As in the literature on the mechanism design approach to patents (for instance Scotchmer [1999]), patents are used to solve information problems; here, however the cost of the patent is serving to get incentives on the planner, rather than provide incentives for the innovator.

[22] For more see Sobel [2011].

The model suggests that philanthropists may be a better way to award subjective prizes, compared to a dispassionate bureaucrat who faces a budget constraint, and therefore may be tempted to not award some prizes that are valuable. Caring about the outcome works as taste, and can therefore lower the information cost of implementing the prize.

7 Conclusion

One interpretation of the differences studied in this paper is as reflecting the *type* of subjectivity. Consider the example of having a building designed. Imagine that the principal wants the building to survive an earthquake of a particular magnitude, and simply does not know how to make a design that can survive that criterion. If the principal's assessment of whether or not the building will survive the earthquake cannot be verified by the courts, it fits into the usual definition of subjective evaluation. The underlying question, however is objective. Subjectivity reflects a lack of perfect judgment on the part of the principal; whether or not the principal can assess value, conditional on the work of the agent, is driven by the relative expertise of the principal and agent. This does not imply that the principal has no judgment; the principal, however, has no extra judgment beyond what the agent's effort offers. In this case noise reflects lack of judgment on the principal's part. In this case contracting can be expensive, and multiple agents can be useful.

Contrast that with the situation where the principal wants a building that looks beautiful to him. The assessment of beauty is inherently subjective: there is no objective test for the principal's opinion of the beauty of the bridge other than his own signal. In this case the principal's signal is a matter of taste and is unambiguously informative about output; beauty is in the eye of the beholder. Therefore the usefulness of right of refusal, and the value of uncertainty when right of refusal is used, is connected with cases where the subjective question is a matter of taste, or where the principal is the expert about the underlying objective question that is being subjectively evaluated.[23] In this case noise reflects not a lack of judgment but rather strong individual preferences. Naturally, then, noise may play a different role in the two cases. This paper elaborates on that role.

[23] Note that an art buyer who wanted to buy a piece that would be widely thought to be beautiful is closer to the subjective evaluation case, since the underlying question "do most people like the way this painting looks?" is objective.

In the model, the project can be refused as a function of the report of the signal. Such an example corresponds to a contract with right of refusal, a common contract form in some types of design. Such a contract is useful to the extent that the principal can forecast his own payoff. This ability to forecast the payoff corresponds naturally to cases of subjective taste, such as in art. Art contracts are one place where right of refusal is common. On the other hand, right of refusal is not useful for cases where the agent is the expert, and therefore can forecast the principal's payoff better than the principal can. Such examples, commonly called credence goods, are often cases where right of refusal is not employed, as the model suggests. Further exploration of the empirical usage of this contract form seems like a useful direction for future work.

More generally the paper defines taste for the principal and explores its value. Although stronger taste, in the sense of the association between signals and the principal's value, is beneficial, having taste correlated with the agent is sometimes not valuable. Moreover, sometimes coarser information is more valuable, meaning that an ignorant patron of the arts might be a more efficient at art procurement.

These relationships held fixed the amount of information or taste that the principal has. An interesting avenue for future work would be to explore the incentive for the principal to invest in taste either ex ante or ex post, and the implications for contracting.

Appendix: Proofs

Proof of Proposition 1

Proof. Suppose $b > 0$. The first order condition for $w(s)$ is

$$\mu_0(\frac{df(s)}{da}/f(s)) + \mu_1 = 0$$

where μ_0 and μ_1 are the Lagrange multipliers on the IC and IR constraints. This can hold for at most one value of s since the MLRP conditions holds. $\square$

Proof of Proposition 2

Proof. Suppose p is strictly increasing for some states s and $s+1$. Then since $p(s)E(s) + w(s) + b(s) \leq p(s+1)E(s) + w(s+1) + b(s+1)$, it must be the case that $p(s)E(s+1)+w(s)+b(s) < p(s+1)E(s+1)+w(s+1)+b(s+1)$, which violates IC. $\square$

Proof of Lemma 3

Proof. For any two s, s' with $p(s) = p(s')$ the same, it must be the case that $w(s) + b(s) = w(s') + b(s')$ is the same, and therefore we can describe the set of distinct contracts, for the purposes of reporting of the principal, by a *strictly* increasing set of distinct contracts $p(s)$.

First we show that, if $1 > p(s) > 0$ and $E(s,a) > a$, some type $s' > s$ must be indifferent between reporting truthfully and reporting s. Suppose therefore that $1 > p(s) > 0$, $E(s,a) > 0$ and all types $s' > s$ strictly prefer truthful reporting to reporting s. Then raise $p(s)$ and $b(s)$ by Δ_p and Δ_b keeping the payoff for s reporting truthfully the same (i.e. let $\Delta_b = \Delta_p E(s,a)$), which make all lower types $s'' < s$ strictly prefer to report truthfully to reporting s. But then nothing binds to s and therefore $p(s) = 1$, a contradiction.

Note that it can never be the case that, for $s' > s > s''$, that both

$$p(s)E(s',a) + w(s) + b(s) = p(s')E(s',a) + w(s') + b(s')$$

and

$$p(s'')E(s',a) + w(s'') + b(s'') = p(s')E(s',a) + w(s') + b(s')$$

(i.e. s' binds to both s and s'') since, if that were the case, then

$$p(s'')E(s,a) + w(s'') + b(s'') > p(s)E(s,a) + w(s) + b(s)$$

i.e. s prefers to report s''. This implies that the contract with the highest $p(s)$ binds to the next highest $p(s)$ (since otherwise nothing would bind to that next highest $p(s)$ from above, as required), and only the next highest $p(s)$ is a candidate to bind to the third-highest $p(s)$, and so on. $\square$

Proof of Proposition 6

Proof. For small enough κ_h, only constraints involving s are impacted. Further, those constraints are both loosened. For reporting type $s+1$:

$$\begin{aligned} p(s)E(s+1,a) + w(s) + b(s) &= p(s+1)E(s+1,a) + w(s+1) + b(s+1) \\ &< p(s+1)\hat{E}(s+1,a) + w(s+1) + b(s+1) \end{aligned}$$

and for reporting type s:

$$\begin{aligned} p(s)E(s,a)+w(s)+b(s) &= p(s-1)E(s,a)+w(s-1)+b(s-1) \\ &> p(s+1)\hat{E}(s-1,a)+w(s+1)+b(s+1) \end{aligned}$$

□

Proof of Proposition 7

Proof. Since $E(1) > 0$, choosing $p_0 = 1$ and $p_1 = 0$ both helps with the reporting constraint and improves the objective. The reporting constraint, if it binds, then says that

$$b_0 = w_1 - w_0 - E(1)$$

So we have that, either $w_1 - w_0 - E(1) < 0$, and $C(\bar{a}) = C^{VS}(\bar{a})$, or the first equation binds and $C(\bar{a}, 0) = C^{VS}(\bar{a}) + (1 - F)(w_1 - w_0 - E(1))$ □

Proof of Proposition 8

Proof. It is optimal to use $p_0 > 0$, rather than b, to satisfy incentives whenever $\lambda > 0$. Moreover $p_1 = 0$ so the IC constraint is

$$w_1 \leq w_0 + p_0 E(1) + b_0$$

If $\omega/E(1) \leq 1$, IC can be satisfied with $b = 0$ and $p_0 = (w_1 - w_0)/E(1)$. When the state is zero this costs $E(0)$ per unit when the state is low, hence the cost in this case is $\omega E(0)/E(1)$.

If $\omega/E(1) > 1$, then $p_0 = 1$ and the principal must burn, in addition to the value of the project $E(0)$, $\omega - E(1)$. □

Proof of Proposition 10

Proof. Since $\lambda > \bar{\lambda}$

$$\frac{d\Gamma}{d\lambda} = -(1-F)\frac{dE(1)}{d\lambda}$$

the cross partial is

$$\frac{d\Gamma}{d\bar{a}d\lambda} = \frac{dE(1)}{d\lambda}\frac{dF}{d\bar{a}} - (1-F)\frac{dE(1)}{d\lambda d\bar{a}}$$

Clearly $E(1)$ is increasing in in λ and F is increasing in $\bar{a}$. The cross partial of $E(1)$ is

$$
\begin{aligned}
\frac{dE(1)}{d\lambda d\bar{a}} &= -\frac{d\frac{\int_{x<\bar{a}} xf(x)dx}{F}}{d\bar{a}} \\
&= -\left[\frac{\bar{a}f(\bar{a})}{F(\bar{a})} - \frac{f(a)\int_{x<\bar{a}} xf(x)dx}{F(\bar{a})^2}\right] \\
&= -\frac{f(\bar{a})}{F(\bar{a})}\left(\bar{a} - E(x|x<\bar{a})\right) < 0
\end{aligned}
$$

Since $\frac{dE(1)}{d\lambda d\bar{a}} < 0$, $\frac{d\Gamma}{d\bar{a}d\lambda} > 0$ □

Proof of Proposition 11

Proof. When $A(0)$ increases, to satisfy IC and IR for the agent, $w(0)$ falls by the same amount.

Suppose $0 < \lambda < \bar{\lambda}$. The fall in $w(0)$ violates the reporting constraint for the principal. If $0 < p(0) < 1$, this is optimally offset by increasing $p(0)$. For each unit of decrease in $w(0)$ the cost of increasing $p(0)$ to maintain the reporting constraint is $E(0)/E(1) < 1$ Therefore C falls. If, on the other hand, $p(a) = 1$, $b(0)$ must increase by the same amount as $w(0)$ falls, leaving C unchanged.

On the other hand suppose $\lambda > \bar{\lambda}$, so $p(0) = 1$. Then if $\Gamma = 0$ the fall in $w(0)$ need not be offset by b and costs fall. On the other hand if $\Gamma > 0$ then, by the same argument as above, $b(0)$ must increase by the same amount as $w(0)$ falls, leaving C unchanged. □

Proof of Lemma 12

Proof. Denoting $F(a+\theta) = F_\theta$ and $E(s,a) = E_\theta(s)$

$$
\begin{aligned}
C_\Phi(a) = &\quad min_{p_i,w_i,l_i,b_i} \int \begin{pmatrix} F_\theta(w_1 + p_1 E_\theta(1) + b_1) + \\ (1-F_\theta)(w_0 + p_0 E_\theta(0) + b_0) \end{pmatrix} d\Phi \\
& \qquad s.t. \\
\int (F_\theta u(w_1) + (1-F_\theta)u(w_0) - c)\, d\Phi \geq &\quad .5u(w_1) + .5u(w_0) \\
0 \leq &\quad \int (F_\theta u(w_1) + (1-F_\theta)u(w_0) - c)\, d\Phi \\
\int (w_1 + p_1 E_\theta(1) + b_1)\, d\Phi \leq &\quad \int (w_0 + p_0 E_\theta(1) + b_0)\, d\Phi \\
\int (w_1 + p_1 E_\theta(0) + b_1)\, d\Phi \geq &\quad \int (w_0 + p_0 E_\theta(0) + b_0)\, d\Phi
\end{aligned}
$$

it is immediate that the constraints reduce to being written as in (3), replacing F with F_Φ and $E(s)$ with $E_\Phi(s)$. Moreover since IR and IC both bind it must be that $w_1 - w_0 = \omega_\Phi$. The objective does not evaluate to the same expression. However:

$\lambda \geq \bar{\lambda}$: The same arguments as in Proposition 7 implies that $p_1 = 0, p_0 = 1$ and therefore $b_1 = max\{\omega_\Phi - E_\Phi(1), 0\}$. Therefore $C_\Phi$is identical to C^{VS} except for b_1, and the expected cost of b_1 is $(1-F_\Phi)b_1$.

$\lambda = 0$: Since $p_0 = p_1 = 0$ is weakly optimal, the argument for $\lambda \geq \bar{\lambda}$ can be applied directly. $C_\Phi = C^{VS} + (1-F_\Phi)\omega_\Phi$. □

Proof of Lemma 13

Proof. Let Φ_u be the unconditional distribution, i.e. before the state is revealed in the first period. Directly computing:

$$\Phi_0(x) = \frac{\int_{\theta<x}(1-F(a+\theta))d\Phi}{\int_\theta (1-F(a+\theta))d\Phi} = \frac{\Phi(x) - \int_{\theta<x} F(a+\theta)d\Phi}{1 - \int_\theta F(a+\theta)d\Phi}$$

and

$$\Phi_1(x) = \frac{\int_{\theta<x} F(a+\theta)d\Phi}{\int_\theta F(a+\theta)d\Phi}$$

Multiplying both by $\frac{1-\int_\theta F(a+\theta)d\Phi}{\int_{\theta<x} F(a+\theta)d\Phi}$ the former is

$$\frac{\Phi(x) - \int_{\theta<x} F(a+\theta)d\Phi}{\int_{\theta<x} F(a+\theta)d\Phi} = \frac{\Phi(x)}{\int_{\theta<x} F(a+\theta)d\Phi} - 1$$

and the latter is

$$\frac{1 - \int_\theta F(a+\theta)d\Phi}{\int_\theta F(a+\theta)d\Phi} = \frac{1}{\int_\theta F(a+\theta)d\Phi} - 1$$

Since

$$\frac{\int_{\theta<x} F(a+\theta)d\Phi}{\Phi(x)} < \int_\theta F(a+\theta)d\Phi$$

it follows that$\Phi_0(x) > \Phi_1(x)$. □

Proof of Proposition 14

Proof. For any a, $V(a, \Phi_s) = a + \int \theta d\Phi_s - C_{\Phi_s}(a)$. F and E are increasing in s; as a result, ω and C^{VS} decline in s. This directly implies that C_{Φ_s} declines in s. Therefore V increases in s. □

References

Richard E Caves. Contracts between art and commerce. *Journal of Economic Perspectives*, pages 73–84, 2003.

Yeon-Koo Che and Ian Gale. Optimal design of research contests. *American Economic Review*, 93(3):646–671, 2003.

P. Crenshaw and R.H. van Rijn. *Rembrandt's Bankruptcy: The Artist, His Patrons, and the Art Market in Seventeenth-Century Netherlands.* Cambridge University Press, 2006. ISBN 9780521858250. URL `https://books.google.ca/books?id=vat8JSgqjVUC`.

Richard Florida. The rise of the creative class and how it's transforming work, leisure, community and everyday life (paperback ed.), 2004.

William Fuchs. Contracting with repeated moral hazard and private evaluations. *American Economic Review*, 97 (4):1432–1448, 2007. doi: 10.1257/aer.97.4.1432. URL `http://www.aeaweb.org/articles.php?doi=10.1257/aer.97.4.1432`.

Suman Ghosh and Michael Waldman. Standard promotion practices versus up-or-out contracts. *The RAND Journal of Economics*, 41(2):301–325, 2010. ISSN 1756-2171. doi: 10.1111/j.1756-2171.2010.00101.x. URL `http://dx.doi.org/10.1111/j.1756-2171.2010.00101.x`.

Oliver Hart and John Moore. Incomplete contracts and renegotiation. *Econometrica: Journal of the Econometric Society*, pages 755–785, 1988.

Bengt Holmstrom. Moral hazard in teams. *The Bell Journal of Economics*, pages 324–340, 1982.

Charles Kahn and Gur Huberman. Two-sided uncertainty and" up-or-out" contracts. *Journal of Labor Economics*, 6(4):423–44, 1988.

Edward P Lazear and Sherwin Rosen. Rank-order tournaments as optimum labor contracts, 1979.

Jonathan Levin. Relational incentive contracts. *The American Economic Review*, 93(3):835–857, 2003.

Nancy A Lutz. Warranties as signals under consumer moral hazard. *Rand Journal of Economics*, 20(2):239–255, 1989.

Bentley W MacLeod. Optimal contracting with subjective evaluation. *The American Economic Review*, 93(1):216–240, 2003.

James M Malcomson. Work incentives, hierarchy, and internal labor markets. *The Journal of Political Economy*, 92(3):486, 1984.

Duncan P Mann and Jennifer P Wissink. Money-back contracts with double moral hazard. *The RAND Journal of Economics*, pages 285–292, 1988.

Suzanne Scotchmer. On the optimality of the patent renewal system. *RAND Journal of Economics*, 30(2):181–196, Summer 1999.

D. Sobel. *Longitude*. HarperCollins Publishers, 2011. ISBN 9780007382002. URL `https://books.google.com/books?id=4Yj8-1xrt6YC`.

Hidenori Takahashi. Strategic design under uncertain evaluations: Theory and evidence from design-build auctions. 2014.

Jan Zabojnik and Dan Bernhardt. Corporate tournaments, human capital acquisition, and the firm size-wage relation. *Review of Economic Studies*, 68(3):693–716, July 2001. URL http://ideas.repec.org/a/bla/restud/v68y2001i3p693-716.html.