

Multi-Observation Elicitation

Sebastian Casalaina-Martin

CU Boulder

CASA@MATH.COLORADO.EDU

Rafael Frongillo

CU Boulder

RAF@COLORADO.EDU

Tom Morgan

Harvard

TDMORGAN@SEAS.HARVARD.EDU

Bo Waggoner

UPenn

BWAG@SEAS.UPENN.EDU

Abstract

We study loss functions that measure the accuracy of a prediction based on multiple data points simultaneously. To our knowledge, such loss functions have not been studied before in the area of property elicitation or in machine learning more broadly. As compared to traditional loss functions that take only a single data point, these multi-observation loss functions can in some cases drastically reduce the dimensionality of the hypothesis required. In elicitation, this corresponds to requiring many fewer reports; in empirical risk minimization, it corresponds to algorithms on a hypothesis space of much smaller dimension. We explore some examples of the tradeoff between dimensionality and number of observations, give some geometric characterizations and intuition for relating loss functions and the properties that they elicit, and discuss some implications for both elicitation and machine-learning contexts.

Keywords: Property elicitation, loss functions, empirical risk minimization.

1. Introduction

In machine learning and statistics, empirical risk minimization (ERM) is a dominant inference technique, wherein a model is chosen which minimizes some loss function over a data set. As the choice of loss function used in ERM may have a large impact on the model chosen, how should one choose this loss? A growing body of work in *property elicitation* seeks to answer this question, by viewing a loss function as “incentivizing” the prediction of a particular conditional statistic (Lambert et al., 2008; Gneiting, 2011; Steinwart et al., 2014; Frongillo and Kash, 2015a; Agarwal and Agarwal, 2015); for example, it is well-known that squared loss elicits the mean, and hence least-squares regression finds the best fit to the conditional means of the data.¹

A natural question, which is still open in the vector-valued case, is the following: for which conditional statistics do there exist loss functions which elicit them? Positive examples include the mean, median, other quantiles, moments, and several others. Perhaps surprisingly, however, there are negative examples as well: it is well-known that the variance is not elicitable, meaning there is no loss function for which minimizing the loss will yield the variance of the data or distribution.

1. There are also contributions from microeconomics, and crowdsourcing in particular, where one wishes to incentivize humans rather than algorithms, but the mathematics is the same.

The usual approach to dealing with non-elicitable statistics is called *indirect elicitation*: elicit other conditional statistics from which one can compute the desired statistic. For example, the variance of a distribution can be written as (2nd moment) - (1st moment)², and as mentioned above, moments are elicitable. The question of how many such auxiliary statistics are required gives rise to the concept of *elicitation complexity*; since the variance cannot be elicited with one but can with two, we say it is 2-elicitable (Lambert et al., 2008; Frongillo and Kash, 2015c).

In this paper, we explore an alternative approach to dealing with non-elicitable statistics, by allowing the loss function to *depend on multiple data points* simultaneously. In the language of property elicitation, this corresponds to loss functions such as $\ell(r, y_1, y_2)$ which judge the “correctness” of the report r based on two (or more) observations y_1 and y_2 . Assuming these observations are drawn independently from the same distribution, this intuitively gives the loss function more power, and could potentially render previously non-elicitable statistics elicitable. In fact, the variance is one such example: if y_1 and y_2 are both drawn i.i.d. from p , it is easy to see that $\frac{1}{2}(y_1 - y_2)^2$ will be an unbiased estimator for the variance of p , hence $\ell(r, y_1, y_2) = (r - \frac{1}{2}(y_1 - y_2)^2)^2$ elicits the variance for the usual reason that squared error elicits expected values. Examples of settings where such i.i.d. observations are readily obtained include: active learning, uncertainty quantification & robust engineering design (Beyer and Sendhoff, 2007), and replication of scientific experiments.

Beyond the variance, are there other non-elicitable statistics which we can elicit with multiple i.i.d. observations? Moreover, what is the tradeoff between the number of observations and the number of reports? One would expect the elicitation complexity, in the usual number-of-reports sense, to drop as observations are added, but how fast is unclear. Indeed, we will see several examples where the complexity drops dramatically, such as the k -norm of the distribution p . In Section 4 we develop new techniques to prove complexity bounds using algebraic geometry, which show for example that the complexity of the k -norm drops from the support size of p (minus 1) with 1 observation, to 1 with k observations. We call the feasible (# reports, # observations) pairs the *elicitation frontier*, for which the given statistic is elicitable, a concept we explore in Section 5.

Finally, in Section 6 we apply multi-observation elicitation to regression. Traditional elicitation complexity expresses a conditional statistic Γ as a link of other statistics, but as we illustrate, situations can arise where these other statistics have a much more complicated relationship with the covariates than Γ does. We give an example where fitting a model to the conditional variance directly (using nearby data points as proxies for i.i.d. observations) is much better than fitting separate models to the conditional first and second moments and combining these to obtain the variance.

1.1. Related work

Our work is inspired in part by Frongillo et al. (2015) which proposes a way to elicit the confidence (inverse of variance) of an agent’s estimate of the bias of a coin by simply flipping it twice. In our terminology, this follows from the fact that the variance is (1, 2)-elicitable. Multi-observation losses have been previously introduced to learn embeddings (Hadsell et al., 2006; Schroff et al., 2015; Ustinova and Lempitsky, 2016), though an explicit property/statistic is never discussed.

2. Preliminaries

We are interested in a space \mathcal{Y} from which *observations* y are drawn, which will be a finite set unless otherwise specified. We will denote by $\mathcal{P} \subseteq \Delta_{\mathcal{Y}}$ a set of probability distributions of interest. (Generally in this paper, \mathcal{P} is simply the entire simplex.) We refer to the set $\Delta_{\mathcal{Y}^m}$ of all distributions

on m outcomes as the m -product space. To capture the assumption that we may collect $m \in \{1, 2, \dots\}$ observations which are each i.i.d. from the same distribution $p \in \Delta_{\mathcal{Y}}$, we will write $p^m \in \Delta_{\mathcal{Y}^m}$ to denote their joint distribution, $p^m(y_1, \dots, y_m) = \prod_i p(y_i)$. The set of all such distributions is denoted $\mathcal{P}^m = \{p^m : p \in \mathcal{P}\} \subseteq \Delta_{\mathcal{Y}^m}$, which we will think of as a manifold in the m -product space.

With this notation in hand, we can define the central concepts in elicitation complexity in our context. Properties include any typical statistic,² for instance, the mean when $\mathcal{Y} \subseteq \mathbb{R}$ is the property $\Gamma(p) = \sum_y p(y)y$.

Definition 1 (Property) A property is a function $\Gamma : \mathcal{P} \rightarrow \mathcal{R}$, where $\mathcal{R} \subseteq \mathbb{R}^k$ for some $k \geq 1$.

Intuitively, properties represent the information desired about the data or underlying distribution. \mathcal{R} is sometimes called the *report space*. The central notion of property elicitation is the relationship between a loss function ℓ and the minimizer of its expected loss. If this minimizer is a particular property Γ , we say ℓ *elicits* Γ . We simply extend this usual definition to allow for multiple observations in the expected loss.

Definition 2 (Loss function, elicits) An m -observation loss function is a function $\ell : \mathcal{R} \times \mathcal{Y}^m \rightarrow \mathbb{R}$, where $\ell(r, y_1, \dots, y_m)$ is the loss for prediction $r \in \mathcal{R}$ scored against realized observations $y_i \in \mathcal{Y}$. We say ℓ (directly) elicits a property $\Gamma : \mathcal{P} \rightarrow \mathcal{R}$ if for all $p \in \mathcal{P}$ we have $\{\Gamma(p)\} = \operatorname{argmin}_{r \in \mathcal{R}} \mathbb{E}_{(y_1, \dots, y_m) \sim p^m} [\ell(r, y_1, \dots, y_m)]$.

It is useful to consider a property in terms of its *level sets*, the set of distributions sharing the same particular value of the property. For example, when the property is the mean of a distribution on $\{1, 2, 3, 4\}$, both $p = (\frac{1}{2}, 0, 0, \frac{1}{2})$ and $p = (0, \frac{1}{2}, \frac{1}{2}, 0)$ lie in the level set $\Gamma_{2.5}$.

Definition 3 (Level set) A level set Γ_r of a property $\Gamma : \mathcal{P} \rightarrow \mathcal{R}$ is, for $r \in \mathcal{R}$, the set of distributions with property r , i.e. $\Gamma_r = \{p \in \mathcal{P} \mid \Gamma(p) = r\}$.

An important technical condition on a property, and one which we will need for the notion of indirect elicitation, is that it be *identifiable*, meaning that its level sets can be described by linear equalities.

Definition 4 (Identifiable) A property $\Gamma : \mathcal{P} \rightarrow \mathcal{R}$, with $\mathcal{R} \subseteq \mathbb{R}^k$, is identifiable with m observations if there exists some $V : \mathcal{R} \times \mathcal{Y}^m \rightarrow \mathbb{R}^k$ such that $\Gamma(p) = r \iff \mathbb{E}_{p^m} [V(r, \mathbf{y})] = 0 \in \mathbb{R}^k$, where $\mathbf{y} = (y_1, \dots, y_m)$ is drawn from p^m . We also say it is m -identifiable.

Identifiability is a geometric restriction on properties that is intuitively similar to continuity of the property (cf. Lambert et al. (2008); Steinwart et al. (2014)). Technically, observe that *differentiable* loss functions generally elicit an identifiable property, as any local optimum should have $\sum_i \frac{\partial}{\partial r_i} \ell(r, \mathbf{y}) = 0$, meaning that the gradient of ℓ itself gives an identification function. Following Frongillo and Kash (2015a), we will often assume that properties are identifiable.

Notice that any property can be “indirectly” elicited by using a proper scoring rule, which elicits the entire distribution, and then computing the property from the distribution. But this requires a report of dimension $|\mathcal{Y}| - 1$, whereas to indirectly elicit the variance of y , for example, requires just

2. As defined, statistics like the median would not be included unless restrictions were placed on \mathcal{P} for them to be single-valued (distributions in general may have multiple medians); we may instead extend our definition to include set-valued statistics, which would not substantially alter our results, and in fact we do lift this restriction in Section 3.1.

two reports, e.g. $r_1 = \mathbb{E}y$ and $r_2 = \mathbb{E}y^2$, along with a “link function” $\psi(\mathbf{r}) = r_2 - r_1^2$. The question of *elicitation complexity*, studied by Lambert et al. (2008) and Frongillo and Kash (2015c), is how many dimensions d are needed to indirectly elicit the property of interest Γ via some elicitable $\hat{\Gamma} : \mathcal{P} \rightarrow \mathbb{R}^d$; one hopes that d is much smaller than $|\mathcal{Y}|$. Here we augment this question by another degree of freedom: how many dimensions d , and observations m , are needed to indirectly elicit Γ ?

Definition 5 (*(d, m) -elicitable*) *A property $\Gamma : \mathcal{P} \rightarrow \mathcal{R}$ is (d, m) -elicitable if there exists a d -dimensional and identifiable property $\hat{\Gamma} : \mathcal{P} \rightarrow \hat{\mathcal{R}}$ where $\hat{\mathcal{R}} \subseteq \mathbb{R}^d$, an m -observation loss function $\ell : \hat{\mathcal{R}} \times \mathcal{Y}^m \rightarrow \mathbb{R}$, and a “link” function $\psi : \hat{\mathcal{R}} \rightarrow \mathcal{R}$, such that*

$$1. \ell \text{ directly elicits } \hat{\Gamma}, \text{ and } \quad 2. \Gamma(p) = \psi(\hat{\Gamma}(p)).$$

The elicitation frontier of Γ is the set of (d, m) such that Γ is (d, m) -elicitable, but neither $(d-1, m)$ -nor $(d, m-1)$ -elicitable.

We may say that a property’s “report complexity” is d if $(d, 1)$ lies on its frontier, and its “observation complexity” is m if $(1, m)$ does.

2.1. Illustrative example

Recall our observation that the variance is not $(1, 1)$ -elicitable, and the “traditional” fix is to utilize $(2, 1)$ -elicitability: minimize a loss function over two dimensions (say first and second moments), mapping the result to the variance via a link function. We observed instead that it is possible to utilize $(1, 2)$ -elicitability: minimize a loss function that takes two observations over a single scalar, the variance itself. Can this tradeoff be more extreme? In particular, are there cases where additional observations drastically decrease the report complexity? Consider the 2-norm of a distribution: $\Gamma(p) = \|p\|_2 = \sqrt{\sum_y p(y)^2}$. We show in Section 5.2 that $\|p\|_2$ has report complexity $|\mathcal{Y}| - 1$ (where \mathcal{Y} is the outcome set) for 1 observation – no single-observation loss function can do better than solving for the entire distribution. However, recall that $\|p\|_2^2 = \sum_y p_y^2 = \Pr\{y_1 = y_2\}$ for two i.i.d. observations y_1, y_2 , or in other words, $\|p\|_2^2 = \mathbb{E}_p \mathbb{1}\{y_1 = y_2\}$. The two-norm is actually elicitable with two observations and a single dimension using e.g. loss function $\ell(r, y_1, y_2) = (r - \mathbb{1}\{y_1 = y_2\})^2$, then simply computing $\|p\|_2 = \sqrt{r}$. In other words, the two-norm’s elicitation frontier on \mathcal{Y} consists of the points $(|\mathcal{Y}| - 1, 1)$ and $(1, 2)$.

The goal for this paper is to investigate the (algebraic-)geometric reasons underpinning why a property might have low or high observation complexity, as well as providing general results and examples based on these ideas. We next introduce the geometric foundations for this investigation.

3. Geometric Fundamentals

The most basic (yet powerful) lower bound in property elicitation says that elicitable properties’ level sets must be convex sets (Lambert et al., 2008). Indeed, this is used to prove the variance is not $(1, 1)$ -elicitable; but the variance *is* elicitable with two observations. The geometry is not “broken” here, but merely lives in a higher-dimensional space. When reasoning about eliciting a property $\Gamma : \mathcal{P} \rightarrow \mathbb{R}$ using m observations, it is often useful to instead think of eliciting the property using a *single random draw from a distribution on m -tuples of outcomes*.

Remark 6 Since \mathcal{P} is isomorphic to \mathcal{P}^m , a property $\Gamma : \mathcal{P} \rightarrow \mathcal{R}$ is directly elicitable with m observations if and only if the induced property $\Gamma^m : \mathcal{P}^m \rightarrow \mathcal{R}$ is directly elicitable with 1 observation. In particular, a sufficient condition for (d, m) -elicitability of Γ is that there exists some $(d, 1)$ -elicitable $\Gamma' : \Delta_{\mathcal{Y}^m} \rightarrow \mathcal{R}$ that coincides with Γ^m on \mathcal{P}^m . One can elicit Γ using the same loss that elicits Γ' , treating the m -tuple of observations as a single draw from the larger space.

This gives us one initial way to demonstrate that a property is elicitable with m observations. For example, the loss function $\ell(r, a, b) = (r - \frac{1}{2}(a - b)^2)^2$ elicits the variance with two observations a, b , but if we consider distributions on all of $\mathcal{Y} \times \mathcal{Y}$, including non-i.i.d. distributions, it actually is still a valid loss function eliciting a property that coincides with the variance when a, b are i.i.d. To see this, just note that it still elicits an expectation: $\sum_{a,b} p'(a, b) \frac{1}{2}(a - b)^2$ where p' is a distribution on \mathbb{R}^2 .

However, considering elicitation on the larger space $\Delta_{\mathcal{Y}^m}$ does not resolve the problem in either the necessary or sufficient directions. First, \mathcal{P}^m is not a convex set for $m > 1$, so conditions on the convexity of level sets do not naturally extend here. An example of this is shown in Figure 1. Second, coming up with an “extended property” may be difficult or non-obvious. For example, it is not so clear whether the above loss function elicits anything *natural* on $\Delta_{\mathcal{Y}^2}$ (it is not the covariance, for instance, which is zero for i.i.d. distributions). More fundamentally, it is not clear whether such extensions should generally exist. (Proving or constructing a counterexample is an interesting open problem.) In general, we hope to be able to accomplish much more by restricting to \mathcal{P}^m because it is only a tiny $|\mathcal{Y}|$ -dimensional manifold in a $|\mathcal{Y}|^m$ -dimensional space.

A tighter sufficient condition is given by [Frongillo and Kash \(2014\)](#), which states that essentially all loss functions eliciting a property on any set, such as \mathcal{P}^m , also elicit some “extension” of that property on the convex hull of that set. So while the higher-dimensional approach is helpful, it does not preclude reasoning about the space \mathcal{P}^m as a manifold inside $\Delta_{\mathcal{Y}^m}$.

Most significantly, \mathcal{P}^m is not a convex space, which makes lower bounds on elicitation complexity nontrivial as well. However, the result of [Frongillo and Kash \(2014\)](#) shows that it suffices to provide lower bounds for elicitation on the convex hull of \mathcal{P}^m , which we will denote $\text{conv}(\mathcal{P}^m)$. Quite naturally then, we explore what leverage we can gain by reasoning about $\text{conv}(\mathcal{P}^m)$.

Theorem 7 *The property $\Gamma : \mathcal{P} \rightarrow \mathbb{R}$ is not directly elicitable with m observations if there exists $r_1, r_2 \in \Gamma(\mathcal{P})$, $p_{1,1}, \dots, p_{1,k_1} \in \Gamma_{r_1}$, $p_{2,1}, \dots, p_{2,k_2} \in \Gamma_{r_2}$, $\lambda_{1,1}, \dots, \lambda_{1,k_1} \in [0, 1]$ and $\lambda_{2,1}, \dots, \lambda_{2,k_2} \in [0, 1]$ such that $r_1 \neq r_2$, $\sum_{i=1}^{k_1} \lambda_{1,i} = 1$, $\sum_{i=1}^{k_2} \lambda_{2,i} = 1$ and*

$$\sum_{i=1}^{k_1} \lambda_{1,i} p_{1,i}^m = \sum_{i=1}^{k_2} \lambda_{2,i} p_{2,i}^m.$$

In other words, a property is not elicitable if there is a convex combination of one of its level sets in the m -product space that equals a convex combination of another one of its level sets in the m -product space.




Figure 1: The two outcome, two observation probability simplex $\Delta_{\mathcal{Y}^2}$ where $\mathcal{Y} = \{0, 1\}$. The arc is the space of i.i.d. distributions $(\Delta_{\mathcal{Y}})^2$.




Figure 2: **Top:** The red dots and blue dots are each a level set of the fourth central moment of a Bernoulli random variable $Y \sim p$. These correspond to the distributions with fourth central moments .07 and .08 respectively. **Bottom:** The curve is $\Delta_{\mathcal{Y}^2}$ projected into \mathbb{R}^2 , and the colored dots are the level sets of the example above projected into this space. The lines demonstrate that there is a point in $\text{conv}(\Delta_{\mathcal{Y}^2})$ that can be written as a convex combination of either of the two level sets.

Theorem 7 allows us to prove for example that the fourth central moment is not directly elicitable with two observations. Consider a Bernoulli random variable $Y \sim p$, then two of the level sets of the fourth central moment $\Gamma(p) = \mathbb{E}_{Y \sim p}[(Y - E_{Y \sim p}[Y])^4]$ are given in Figure 2. When we project these level sets into the 2-product space we can easily find a pair of points from each level set whose connecting lines intersect in $\text{conv}(\Delta_{\mathcal{Y}^m})$. These lines are convex combinations of points in the same level set, so by Theorem 7 the lines’ intersection implies that Γ is not directly elicitable with two observations.

3.1. Finite Properties

Finite properties are those where \mathcal{R} , the range of Γ , is a finite set. This corresponds to a “multiple-choice question” (Lambert and Shoham, 2009). In this section, we must allow $\Gamma : \mathcal{P} \rightrightarrows \mathcal{R}$ to be a set-valued function, possibly assigning multiple possible correct reports to a single distribution; this is necessary for “boundary” cases, such as the mode of the uniform distribution on a finite set. (Similarly, we cannot require identifiability.) We have a finite set of outcomes \mathcal{Y} , the distributions considered are all $\mathcal{P} = \Delta_{\mathcal{Y}}$, and $\Gamma(p)$ must be nonempty.

We are interested in understanding which finite properties can be elicited with m observations. Previously, this question was studied for the case of one observation by Lambert (2011), who characterized elicitable properties by the shape of their level sets: they are intersections of *Voronoi diagrams* in $\mathbb{R}^{|\mathcal{Y}|}$ with the simplex $\Delta_{\mathcal{Y}}$. In our setting, a Voronoi diagram is specified by a finite set of points $\{x_r : r \in \mathcal{R}\} \subseteq \mathbb{R}^{\mathcal{Y}^m}$, with each *cell* $T_r = \{x : \|x - x_r\| \leq \|x - x_{r'}\| \forall r' \in \mathcal{R}\}$ consisting of those points in $\mathbb{R}^{\mathcal{Y}^m}$ closest in Euclidean distance to x_r .

Using the geometric constructions above, we can simply apply the main result of Lambert (2011) to finite properties in the m -product space; the result is a characterization of elicitable finite properties with m observations.

Corollary 8 *A finite property $\Gamma : \Delta_{\mathcal{Y}} \rightrightarrows \mathcal{R}$ is directly elicitable with m samples if and only if there exists a Voronoi diagram in $\mathbb{R}^{\mathcal{Y}^m}$ with $\{x_r : r \in \mathcal{R}\}$ satisfying $\Gamma_r^m = T_r \cap \mathcal{P}^m$. Here $\Gamma_r^m = \{p^m \in \mathcal{P}^m : p \in \Gamma_r\}$.*

Multiple observations afford considerable flexibility in the level sets of such an elicitable Γ . In particular, whereas before the cell boundaries between level sets were restricted to hyperplanes,

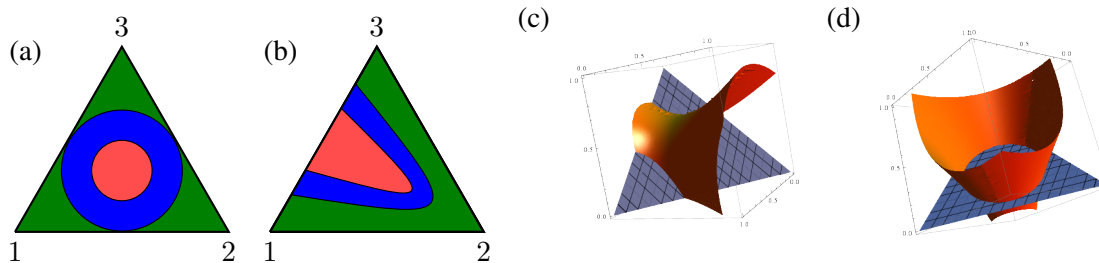


Figure 3: **Examples of finite properties on $\mathcal{Y} = \{1, 2, 3\}$ elicitable with 2 samples.** Pictured is the simplex on 3 outcomes and properties $\Gamma : \Delta_{\{1,2,3\}} \rightarrow \{\text{red, green, blue}\}$. The agent reports a color, then is rewarded according to which outcome occurs. (a) The property of “close”, “intermediate”, and “far” from uniform, as measured by 2-norm. (b) The property of “high”, “medium”, and “low” variance. (c,d) The boundary between two cells, i.e. a hyperplane in the m -product space “projected” down to $\mathbb{R}^{\mathcal{Y}}$ (in orange) and intersected with the simplex $\Delta_{\mathcal{Y}}$ (in blue/gray); we show the boundary on all of $\mathbb{R}^{\mathcal{Y}}$ to visualize the quadratic surfaces which create these sections.

with m observations these boundaries can be defined by nearly arbitrary m -degree polynomials. We illustrate this flexibility and visualize the cell boundaries in Figure 3. In particular, we show that a classic negative example, where an agent is asked to report whether their belief has low or high variance, is easily elicited with two observations.

4. Lower Bounds via Geometry

In this section we discuss lower bounds on elicitation complexity. For technical reasons we will here require \mathcal{P} to be a C^∞ submanifold of $\Delta_{\mathcal{Y}}$ with corners. Our lower bounds will also generally require Γ to be a C^∞ function, in which case we call it a C^∞ property.

We begin in the first subsection by recalling the structure of the level sets of identifiable properties, and then introduce a technique for obtaining from this some lower bounds on elicitation complexity via differential geometry. In the next subsection we focus on polynomial properties, and explain some results that use algebraic geometry to obtain sharp bounds.

4.1. Preliminaries on identifiable properties

We start by recalling a general method, introduced in [Frongillo and Kash \(2015c\)](#), for showing lower bounds on elicitation complexity: *Given a property Γ , if one can show that no level set from any $\hat{\Gamma}$, which is m -identifiable and directly elicitable with m observations, can be contained in a particular level set of Γ , then Γ cannot be (d, m) -elicitable.* This follows immediately from the definitions: if Γ is indirectly elicited via $\hat{\Gamma}$ and link ψ , so that $\Gamma = \psi \circ \hat{\Gamma}$, then we have the following relationship between the level sets of Γ and $\hat{\Gamma}$:

$$\Gamma_r = \bigcup_{\hat{r}:\psi(\hat{r})=r} \hat{\Gamma}_{\hat{r}}. \quad (1)$$

In other words, the level sets of Γ are obtained by combining some of the level sets of $\hat{\Gamma}$. For instance, if ψ is a bijection, then the level sets of Γ and $\hat{\Gamma}$ are identical. This method was used successfully in [Frongillo and Kash \(2015c\)](#) to show lower bounds on the report complexity (d) of

a property, with $m = 1$. In this section, we will use the same method to show lower bounds on observation complexity (m), with $d = 1$.

Our main tool for obtaining these lower bounds will be that the level sets of any directly m -observation-elicitable, identifiable $\hat{\Gamma}$ have a specific structure, namely, such a level set is the zero set of a polynomial of degree at most m :

Fact 1 *If a property $\hat{\Gamma}(p)$ is m -identifiable, then each level set of $\hat{\Gamma}$ is the set of zeros of a polynomial in p of degree at most m .*

Proof The condition $\mathbb{E}_p V(r, \mathbf{y}) = 0$ is $\sum_{y_1, \dots, y_m} p(y_1) \cdots p(y_m) V(r, y_1, \dots, y_m) = 0$. ■

Combined with the equality (1) above, Fact 1 tells us that the level sets of indirectly elicitable Γ are unions of zero sets of polynomials. As we are focusing on the $d = 1$ case, however, both Γ and $\hat{\Gamma}$ are real-valued functions, so with enough regularity, their level sets should coincide. Before making a precise statement, we introduce the following definition:

Definition 9 ($C^\infty(d, m)$ -**elicitable**) *We say that a C^∞ property $\Gamma : \mathcal{P} \rightarrow \mathbb{R}^d$ is $C^\infty(d, m)$ -elicitable if in the definition of (d, m) -elicitable, $\hat{\Gamma}$ can be taken to be C^∞ and ψ can be taken to be C^∞ in an open neighborhood of the image of $\hat{\Gamma}$.*

Corollary 10 *Suppose that a C^∞ property $\Gamma : \mathcal{P} \rightarrow \mathbb{R}$ is $C^\infty(1, m)$ -elicitable. Let $r \in \mathbb{R}$, let $Z \subseteq \Gamma^{-1}(r)$ be a connected component of the level set, and assume that Z admits a point that is not a critical point of Γ ; i.e., there is a point $p \in Z$ such that the differential of Γ at p is nonzero. Then Z is a connected component of the set of zeros of a polynomial of degree at most m . Moreover, if $\Gamma^{-1}(r)$ is connected, then $\Gamma^{-1}(r)$ is the zero set of a polynomial of degree at most m .*

Proof Let $\hat{\Gamma}$ and ψ be as in Definition 9. We have a commutative diagram:

$$\begin{array}{ccc} \mathcal{P} & \xrightarrow{\hat{\Gamma}} & \mathbb{R} \\ \Gamma \downarrow & \searrow \psi & \\ \mathbb{R} & & \end{array}$$

Since Z is connected, we have that $\hat{\Gamma}(Z) \subseteq \mathbb{R}$ is connected, and is therefore an interval (see e.g., Browder (1996), Theorem 6.76, 6.77, p.148). The claim is that this interval is a point. Indeed, assume the opposite. Then since ψ is by definition constant on the interval $\hat{\Gamma}(Z)$, we would have that the differential $D\psi$ vanishes at each point of $\hat{\Gamma}(Z)$. Then since $D\Gamma = D\psi \circ D\hat{\Gamma}$ we would have that $D\Gamma$ vanishes at every point of Z . But this would contradict our assumption. Thus $\hat{\Gamma}(Z)$ is a point.

It then follows from Fact 1 that $\hat{\Gamma}^{-1}(\hat{\Gamma}(Z))$ is the zero set of a polynomial of degree at most m . We now use the inclusions

$$Z \subseteq \hat{\Gamma}^{-1}(\hat{\Gamma}(Z)) \subseteq \Gamma^{-1}(r).$$

By virtue of the inclusion on the right, every connected component of $\hat{\Gamma}^{-1}(\hat{\Gamma}(Z))$ is contained in a connected component of $\Gamma^{-1}(r)$. This proves the first assertion of the lemma. The last assertion of the lemma also follows from these inclusions, since in that case one is assuming $Z = \Gamma^{-1}(r)$. ■

Remark 11 *For concreteness, we summarize the contrapositive of Corollary 10 in the way in which we will use it in examples: Suppose that $\Gamma : \mathcal{P} \rightarrow \mathbb{R}$ is a C^∞ property, and there exists an $r \in \mathbb{R}$ such that the level set $\Gamma^{-1}(r)$ is connected, and contains a point $P \in \Gamma^{-1}(r)$ that is not a critical point for Γ . Then if $\Gamma^{-1}(r)$ is not the zero locus of a degree m polynomial in $p(y_1), \dots, p(y_m)$, then Γ is not $C^\infty(1, m)$ -elicitable.*

As a consequence of Corollary 10, we can immediately show the existence of C^∞ properties with *infinite* observation complexity; i.e., properties that are not $C^\infty(1, m)$ elicitable for any m . The proof gives such an example for $|\mathcal{Y}| = 3$, a surprising result given that all properties have report complexity $|\mathcal{Y}| - 1 = 2$; i.e., all of the C^∞ properties are $C^\infty(2, 1)$ -elicitable. Note that if $|\mathcal{Y}| = 2$, then all C^∞ properties are $C^\infty(1, 1)$ -elicitable.

Proposition 12 *There are C^∞ properties that are not $C^\infty(1, m)$ -elicitable for any finite m .*

Proof Take $\mathcal{Y} = \{1, 2, 3\}$, $\mathcal{P} = \Delta_{\mathcal{Y}}^\circ = \{p \in \Delta_{\mathcal{Y}} : p(y) > 0 \ \forall y \in \mathcal{Y}\}$, and $\Gamma(p) = p_1 - (1/2) \sin(1/p_2)$. It is immediate that Γ has no critical points. Here the level sets Γ_r satisfy $r = p_1 - (1/2) \sin(1/p_2)$, in other words, satisfy the equation $p_1 = (1/2) \sin(1/p_2) + r$. For p_2 sufficiently small, the level set $\Gamma_0 = \{p \in \Delta_3 : p_1 = (1/2) \sin(1/p_2)\}$ is simply the graph of $(1/2) \sin(1/x)$, which intersects the line $p_1 = 0$ infinitely many times, and hence by the Fundamental Theorem of Algebra is not the zero set of *any* polynomial. Corollary 10 now implies that Γ is not $(1, m)$ -elicitable for any m . ■

4.2. Polynomial properties and lower bounds using algebraic geometry

We now describe some lower bounds for elicitation complexity of polynomial properties. The motivation for these lower bounds is the intuition that, in general, a polynomial property $\Gamma : \mathcal{P} \rightarrow \mathbb{R}$ of degree k should not be $C^\infty(1, m)$ -identifiable for any $m < k$, since the zero set of a degree k polynomial should not be the zero set of a degree m polynomial when $m < k$. This statement can of course fail in special cases (e.g., Example 1 below). Indeed, there are some subtleties regarding zero sets of polynomials in Euclidean open sets, considered in Appendix C, that must be addressed to draw such a conclusion. Nevertheless, for a general polynomial property this expectation holds (see Remark C.4 for a precise definition of generality), and in the appendix we provide some elementary techniques for confirming this expectation in particular examples (see Corollary C.3). For instance, we show (Example C.1):

Corollary 13 *If $|\mathcal{Y}| \geq 3$, then for any natural number k , the k -norm of a distribution, $\Gamma(p) = (\sum_y p(y)^k)^{1/k}$, is not $C^\infty(1, k-1)$ -elicitable.*

Example 1 *In contrast to the case considered in Corollary 13, we emphasize that there are polynomial properties $\Gamma : \mathcal{P} \rightarrow \mathbb{R}$ of degree k that are $C^\infty(1, m)$ -elicitable for some $m < k$. For instance, take $\hat{\Gamma} : \mathcal{P} \rightarrow \mathbb{R}$ to be any polynomial property of degree $m > 0$, let $\psi : \mathbb{R} \rightarrow \mathbb{R}$ be any polynomial function of degree $m' > 1$, and set $\Gamma = \psi \circ \hat{\Gamma}$. Then Γ is of degree $k = mm' > m$, but Γ is $C^\infty(1, m)$ -elicitable, by Lemma 15.*

5. Examples and Elicitation Frontiers

We now combine our complexity lower bounds with upper bounds to make progress toward determining the elicitation frontiers of some potential properties of interest. See Figure 4 for a depiction of some of the elicitation frontiers described. We begin with some general, straightforward, but versatile upper bounds.

Lemma 14 *For all $1 \leq i \leq n$, $1 \leq j \leq m$, let $f_{ij} : \mathcal{Y} \rightarrow \mathbb{R}$ be an arbitrary function such that $\mathbb{E}_p[f_{ij}(Y)]$ exists for all $p \in \mathcal{P}$. Then $\Gamma(p) = \sum_{i=1}^n \prod_{j=1}^m \mathbb{E}_p[f_{ij}(Y)]$ is $(1, m)$ -elicitable.*

Proof Using Y_1, \dots, Y_m which are i.i.d. from p , then $\{f_{i1}(Y_1), \dots, f_{im}(Y_m)\}$ will be independent for all i . Using properties of expectations (linearity and independence), we have

$$\sum_{i=1}^n \prod_{j=1}^m \mathbb{E}[f_{ij}(Y)] = \sum_{i=1}^n \prod_{j=1}^m \mathbb{E}[f_{ij}(Y_j)] = \sum_{i=1}^n \mathbb{E} \left[\prod_{j=1}^m f_{ij}(Y_j) \right] = \mathbb{E} \left[\sum_{i=1}^n \prod_{j=1}^m f_{ij}(Y_j) \right] \quad (2)$$

Now we see that using squared loss (or any loss for the mean) one can leverage these m samples to elicit the desired sum of products, e.g. $\ell(r, y_1, \dots, y_m) = \left(r - \sum_{i=1}^n \prod_{j=1}^m f_{ij}(y_j) \right)^2$. ■

The proof of Lemma 14 simply constructs an unbiased estimator of the property of interest and elicits the mean of the estimator via squared error. By a very natural extension, this technique also applies to ratios of expectations, as they are elicitable (Gneiting, 2011): construct *two* unbiased estimators, and elicit the ratio of their means. We will give two instances of such ratios in the next subsection.

The following result establishes an upper bound that by now may seem natural: Under some conditions, a property that is itself an m -degree polynomial in p is $(1, m)$ -elicitable.

Lemma 15 *Suppose that $\Gamma : \mathcal{P} \rightarrow \mathbb{R}$ is a property such that $\Gamma = \psi \circ \Gamma'$ where $\Gamma' : \mathcal{P} \rightarrow \mathbb{R}$ is polynomial of degree m , and $\psi : \mathbb{R} \rightarrow \mathbb{R}$ is a function that is C^∞ on an open neighborhood of the image of Γ' . Then Γ' is directly $(1, m)$ -elicitable, and Γ is C^∞ $(1, m)$ -elicitable.*

Proof It is enough to show Γ' is directly $(1, m)$ -elicitable. This follows immediately from Lemma 14. Indeed, it is clear from the lemma that it is enough to show the result for monic monomials. For this one takes the f_{ij} in Lemma 14 to be characteristic functions $\mathbb{1}_y$ for $y \in \mathcal{Y}$. ■

5.1. Ratios of expectations: index of dispersion and Sharpe ratio

The *index of dispersion* of a random variable Y with positive mean is defined to be $\text{Var}(Y)/\mathbb{E}[Y]$ (Cox and Lewis, 1966). The *Sharpe ratio* of a random variable Y , which is a commonly-used measure of the risk-adjusted return of an investment, is defined similarly as $\mathbb{E}[Y]/\sqrt{\text{Var}(Y)}$ (Sharpe, 1966). Both the index of dispersion and the *square* of the Sharpe ratio are $(1, 2)$ -elicitable by the above discussion: $\text{Var}(Y) = \mathbb{E}_p[\frac{1}{2}(Y_1 - Y_2)^2]$, $\mathbb{E}_p[Y] = \mathbb{E}_p[Y_1]$, and $\mathbb{E}_p[Y]^2 = \mathbb{E}_p[Y_1 Y_2]$, so any ratios of these terms is $(1, 2)$ -elicitable. (The link function for the Sharpe ratio is thus the square root.) For example, the index of dispersion is elicited by the loss $\ell(r, y_1, y_2) = r(y_1 - y_2)^2 - r^2 y_1$.

To finish describing the elicitation frontiers for these properties, we note that neither is $(1, 1)$ -elicitable as the level sets are not convex, but both are $(2, 1)$ -elicitable as we now show. For the

index of dispersion, we can take $r_1 = E[Y]$ and $r_2 = E[Y^2]$, both elicitable as means, and then compute the property by $(r_2 - r_1^2)/r_1$. Similarly, for the same r_1, r_2 , the Sharpe ratio can be written as $r_1/\sqrt{r_2 - r_1^2}$.

5.2. Norms of distributions

As we have previously discussed, the 2-norm is $(1, 2)$ elicitable. For general k , the k -norm is $(1, k)$ elicitable with the following loss function $\ell(r, y_1, \dots, y_k) = (r - \mathbb{1}\{y_1 = \dots = y_k\})^2$. (This case also follows from Lemma 15.) This is a tight bound on the observation complexity, as we proved in Corollary 13 that the k -norm is not $(1, k-1)$ elicitable. As it turns out, the report complexity of the k -norm is $|\mathcal{Y}| - 1$, meaning it is as hard to elicit with one observation as the entire distribution. This follows from Theorem 2 of Frongillo and Kash (2015c), specifically Section 4.2, as $\|p\|_k$ is a convex function of p . An interesting open question, and one that will require additional algebraic tools, is the k -norm's elicitation frontier when we allow multiple dimensions and multiple observations.

Corollary 16 For $|\mathcal{Y}| \geq 3$, the elicitation frontier of the k -norm contains $(|\mathcal{Y}| - 1, 1)$ and $(1, k)$.

5.3. Central Moments

The n^{th} central moment μ_n of a random variable Y is defined as

$$\mu_n = \mathbb{E}[(Y - \mathbb{E}[Y])^n] = \sum_{i=0}^n (-1)^i \binom{n}{i} \mathbb{E}[Y]^i \cdot \mathbb{E}[Y^{n-i}], \quad (3)$$

which we see is $(n, 1)$ -elicitable by simply eliciting $E[Y^i]$ for all $i \in \{1, \dots, n\}$ and then combining the results. As we will show, μ_n is also $(1, n)$ -elicitable, and moreover, we can achieve other dimension-observation tradeoffs in between, such as $(\lfloor \sqrt{n} \rfloor + 1, \lceil \sqrt{n} \rceil)$. The key idea is to partition the binomial sum (3) into k partial sums and factor out the highest power of $\mathbb{E}[Y]$ from each, such that the j^{th} partial sum can be written as

$$\mathbb{E}[Y]^{j \cdot \frac{n}{k}} \sum_{i=0}^{\frac{n}{k} - 1} (-1)^i \binom{n}{j \cdot \frac{n}{k} + i} \cdot \mathbb{E}[Y]^i \cdot \mathbb{E}\left[Y^{\frac{(j+1) \cdot n}{k} - 1 - i}\right]. \quad (4)$$

Doing so gives the following result.

Theorem 17 The n^{th} central moment is $(k + 1, \lceil n/k \rceil)$ -elicitable; $0 < k \leq n$

Proof Consider the partial sum (4) without the $\mathbb{E}[Y]^{j \cdot n/k}$ factor; by Lemma 14 each such factored sum is $(1, \lceil n/k \rceil)$ -elicitable, as the maximum number of terms in any product is $\lceil n/k \rceil$. Since we have k such factored sums, and need to additionally elicit the mean $\mathbb{E}[Y]$ to compute their factors, the entire sum can be elicited using $\lceil n/k \rceil$ observations and $k + 1$ dimensions. ■

When $k = 0$, we can do much better than $m = \infty$: by Lemma 14, as the maximum number of terms in any product of (3) is n , the term $(\mathbb{E}[Y])^n$, we have that μ_n is $(1, n)$ -elicitable. For lower bounds, little is known beyond μ_n not being $(1, 1)$ -elicitable (Frongillo and Kash, 2015b).

6. Multi-Observation Regression

One of the earliest problems in modern statistics was the estimation of biodiversity in a geographic region (Fisher et al., 1943). One scalar measure of diversity of a distribution is the (inverse of the) 2-norm, which we will take here as an example.³ Consider a dataset of species samples: pairs (x, y) where x gives the features of the geographic region and y is a categorical giving the species to which this sample belongs. Suppose we wish to regress the diversity of species against geographic features such as climate. The single-observation approach would require a surrogate loss function $\ell(\hat{f}(x), y)$ and a link $f(x) = \psi(\hat{f}(x))$. We claim that any single-observation loss function $\ell(f(x), y)$ is poorly suited for this task. For the 2-norm, lower bounds on report complexity show that the best possible approach has dimensionality $\hat{f} : x \rightarrow \mathbb{R}^{d-1}$ where d is the number of unique species in the dataset (which may have a very long tail). So this approach requires, in essence, fitting \hat{f} to the entire distribution over species as a function of geographic region, a task of immense idiosyncrasy and complexity compared to the end goal of e.g. estimating a scalar measure of diversity as a function of rainfall level.

On the other hand, a two-observation loss function $\ell(f(x), y_1, y_2)$ can be used to directly learn an f estimating the desired diversity measure, e.g. 2-norm, as a function of geographic features. One can then use empirical risk minimization to directly learn relationships between, e.g. rainfall level and this measure of species diversity.

Multi-observation regression does introduce an additional challenge, however: risk in this context is naturally defined as $\mathbb{E}_{x, \mathbf{y}} \ell(f(x), \mathbf{y})$ where $\mathbf{y} = (y_1, \dots, y_m)$ is a set of observations drawn i.i.d. conditioned on x , but our data points are of the form (x, y) . If e.g. x comes from a continuous space, we may not have *any* sets of m samples y_1, \dots, y_m belonging to the same x . One natural setting where this poses no concern is in active learning where we may choose to re-draw the label for a given x . In a more standard regression framework, we propose to leverage the intuition that the distribution of y conditioned on x generally changes gradually as a function of x .⁴ Pragmatically, with dense enough data points, we can simply group together nearby x values and “merge” them into a data point of the form $(\bar{x}, y_1, \dots, y_m)$ where \bar{x} is an average and the y_i are drawn independently and *approximately* identically from *approximately* the distribution of \mathcal{Y} conditioned on \bar{x} . For this paper, we demonstrate the idea in simulations below and give a basic proof-of-concept theoretical result in Appendix B, leaving a more thorough investigation to future work.

3. A similar intuition will hold for most if not all elicitable measures of diversity.

4. Phrased differently, at least it seems reasonable to parameterize the rate of change and expect learning bounds to depend on this parameter.




Figure 4: The elicitation frontiers for various properties: (a) the variance, Sharpe ratio, and index of dispersion; (b) the 2-norm when $|\mathcal{Y}| = 6$, with respect to C_∞ properties; (c) $\Gamma(p) = p_1 - (1/2) \sin(p_2\pi)$ from Proposition 12; (d) the 4th central moment, which is not fully known.

In general, the cases where the multi-observation approach can be useful are those where the property of interest is believed to follow a simple functional form, but the conditional statistics given by the indirect elicitation approach are expected to follow unknown or complicated trends as a function of features. For another example, one could imagine learning the noise (e.g. variance) of a medical test, e.g. white blood cell count, as a function of patient features, in order to improve the test. The indirect elicitation approach suggests first fitting a model for estimating the mean of the test’s outcome as a function of patient data, then fitting the expected square of the statistic, and then computing an estimate for the variance by combining them. In general, these prediction problems may be highly complex and nonlinear even when the *noise* in the test might follow some simple linear relationship with e.g. height or age. The multi-observation approach allows direct regression of the noise versus features. Formally, we show a basic extension of classic risk guarantees in Appendix B, under the assumption that x is distributed uniformly on $[0, 1]$ and a closeness condition on the conditional distribution of Y given X .

6.1. Simulation

Here we describe some simulations run as a proof of concept of multi-observation regression. Our data points are of the form $(x, y) \in \mathbb{R} \times \mathbb{R}$ where x is drawn uniformly at random from the interval $[0, 1]$. Given x , $y = a \sin(4\pi x) + Z$, where a is a constant and $Z \sim N(0, 1)$ is drawn independently for each sample, we wish to learn $\text{Var}(Y|X)$.

Our multi-observation loss function here is $\ell(f(x), y_1, y_2) = (f(x) - \frac{1}{2}(y_1 - y_2))^2$. We approximate (x, y_1, y_2) samples by sorting the (x_i, y_i) pairs by x_i , and making samples of the form $(\frac{1}{2}(x_i + x_{i+1}), y_i, y_{i+1})$. We compare to the single observation approach, in which we estimate $\mathbb{E}[Y|X]$ and $\mathbb{E}[Y^2|X]$ and then combine them to estimate $\text{Var}(Y|X)$.

The point of these simulations is to demonstrate that multi-observation regression can greatly outperform single observation regression in the case when the function is in a known concept class, and the statistics needed to indirectly elicit it with a single observation are not in a known concept class. As such, our multi-observation regression fits a linear function to $\text{Var}(Y|X)$, and our single observation regression fits linear functions to $\mathbb{E}[Y|X]$ and $\mathbb{E}[Y^2|X]$. The true $\text{Var}(Y|X) = 1$ is indeed a linear function, while the true moment functions $\mathbb{E}[Y|X = x] = a \sin(x)$ and $\mathbb{E}[Y^2|X = x] = a^2 \sin^2(x) + 1$ are very far from linear.

Figure 5 gives the results for $a = 1$ and $a = 10$. Both plots show the mean squared error of the variance functions reported by the two regression methods (averaged over 4000 simulations) as a function of the number of samples. In both cases we see that for sufficiently many samples, the two observation regression significantly outperforms the single observation regression.

7. Conclusion and Future Work

An immediate host of directions is the proving of upper and lower bounds on elicitation frontiers for various properties. In particular, our lower bounds here focus on techniques for lower-bounding observation complexity (the $(1, m)$ case), leaving open approaches for lower bounds on (d, m) complexity for $d \geq 2$. Another direction is to formalize learning guarantees for multi-observation regression under suitable assumptions on slow-changing conditional distributions.

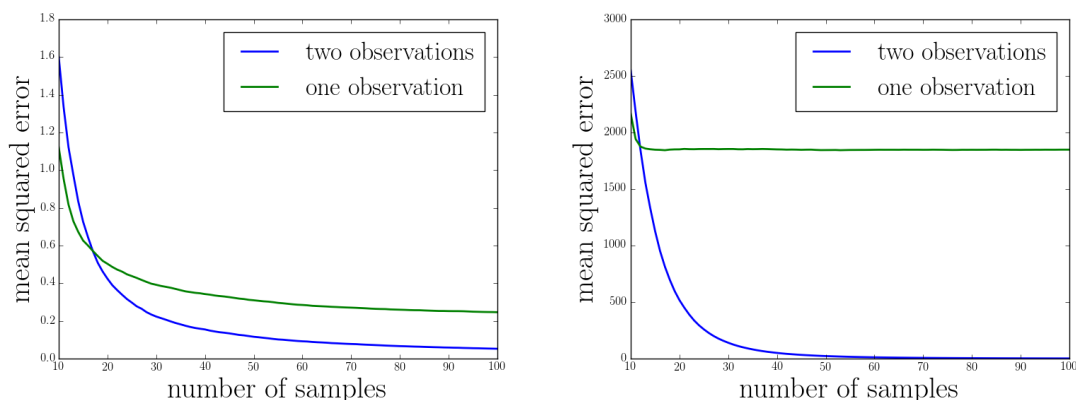


Figure 5: The mean squared error of the two regression strategies for estimating $\text{Var}(y|x)$, where $x \sim \text{Unif}(0, 1)$ and $y \sim a \sin(4\pi x) + N(0, 1)$, for $a = 1$ (left) and $a = 10$ (right). The single-observation loss function approach fails because it tries to fit to the complex underlying model of $y|x$, while the two-observation loss approach is able to directly model the simple relationship between $\text{Var}(y)$ and x .

Acknowledgments

We thank Karthik Kannan for contributing the upper bound for central moments. Sebastian Casalaina-Martin was partially supported by NSA grant H98230-16-1-0053. Tom Morgan was funded in part by NSF grants CCF-1320231 and CNS-1228598. Bo Waggoner is supported by the Warren Center for Network and Data Sciences at the University of Pennsylvania.

References

- Arpit Agarwal and Shivani Agarwal. On consistent surrogate risk minimization and property elicitation. In *JMLR Workshop and Conference Proceedings*, volume 40, pages 1–19, 2015.
- M. F. Atiyah and I. G. Macdonald. *Introduction to commutative algebra*. Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont., 1969.
- Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- Hans-Georg Beyer and Bernhard Sendhoff. Robust optimization—a comprehensive survey. *Computer methods in applied mechanics and engineering*, 196(33):3190–3218, 2007.
- Jacek Bochnak, Michel Coste, and Marie-Françoise Roy. *Real algebraic geometry*, volume 36 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1998. ISBN 3-540-64663-9. Translated from the 1987 French original, Revised by the authors.
- Andrew Browder. *Mathematical analysis: An introduction*. Undergraduate Texts in Mathematics. Springer-Verlag, New York, 1996. ISBN 0-387-94614-4.

- David A. Cox, John Little, and Donal O’Shea. *Ideals, varieties, and algorithms*. Undergraduate Texts in Mathematics. Springer, Cham, fourth edition, 2015. ISBN 978-3-319-16720-6; 978-3-319-16721-3. An introduction to computational algebraic geometry and commutative algebra.
- David R Cox and Peter AW Lewis. The statistical analysis of series of events. *Monographs on Applied Probability and Statistics*, 1966.
- Ronald A. Fisher, A. Steven Corbet, and Carrington B. Williams. The relation between the number of species and the number of individuals in a random sample of an animal population. *The Journal of Animal Ecology*, 12(1):42–58, 1943.
- Rafael Frongillo and Ian Kash. General truthfulness characterizations via convex analysis. In *Web and Internet Economics*, pages 354–370. Springer, 2014.
- Rafael Frongillo and Ian Kash. Vector-Valued Property Elicitation. In *Proceedings of the 28th Conference on Learning Theory*, pages 1–18, 2015a.
- Rafael Frongillo and Ian A. Kash. On Elicitation Complexity and Conditional Elicitation. *arXiv preprint arXiv:1506.07212*, 2015b.
- Rafael Frongillo and Ian A. Kash. On Elicitation Complexity. In *Advances in Neural Information Processing Systems 29*, 2015c.
- Rafael M. Frongillo, Yiling Chen, and Ian A. Kash. Elicitation for Aggregation. *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, 2015.
- William Fulton. *Intersection theory*, volume 2 of *Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics [Results in Mathematics and Related Areas. 3rd Series. A Series of Modern Surveys in Mathematics]*. Springer-Verlag, Berlin, second edition, 1998. ISBN 3-540-62046-X; 0-387-98549-2.
- T. Gneiting. Making and Evaluating Point Forecasts. *Journal of the American Statistical Association*, 106(494):746–762, 2011.
- Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 1735–1742. IEEE, 2006.
- Nicolas S. Lambert and Yoav Shoham. Eliciting truthful answers to multiple-choice questions. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, pages 109–118, 2009.
- Nicolas S. Lambert, David M. Pennock, and Yoav Shoham. Eliciting properties of probability distributions. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 129–138, 2008.
- N.S. Lambert. Elicitation and Evaluation of Statistical Forecasts. *Preprint*, 2011.
- Kent Harold Osband. *Providing Incentives for Better Cost Forecasting*. University of California, Berkeley, 1985.

Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 815–823, 2015.

William F Sharpe. Mutual fund performance. *Journal of Business*, 39(1):119–138, 1966.

Ingo Steinwart, Chloé Pasin, Robert Williamson, and Siyu Zhang. Elicitation and Identification of Properties. In *Proceedings of The 27th Conference on Learning Theory*, pages 482–526, 2014.

Evgeniya Ustinova and Victor Lempitsky. Learning deep embeddings with histogram loss. In *Advances in Neural Information Processing Systems*, pages 4170–4178, 2016.