
Leveraging Contextual Counterfactuals Toward Belief Calibration

Qiuyi (Richard) Zhang¹ Michael S. Lee² Sherol Chen¹

Abstract

Beliefs and values are increasingly being incorporated into our AI systems through alignment processes, such as carefully curating data collection principles or regularizing the loss function used for training. However, the meta-alignment problem is that these human beliefs are diverse and not aligned across populations; furthermore, the implicit strength of each belief may not be well calibrated even among humans, especially when trying to generalize across contexts. Specifically, in high regret situations, we observe that contextual counterfactuals and recourse costs are particularly important in updating a decision maker’s beliefs and the strengths to which such beliefs are held. Therefore, we argue that including counterfactuals is key to an accurate calibration of beliefs during alignment. To do this, we first segment belief diversity into two categories: subjectivity (across individuals within a population) and epistemic uncertainty (within an individual across different contexts). By leveraging our notion of epistemic uncertainty, we introduce ‘the belief calibration cycle’ framework to more holistically calibrate this diversity of beliefs with context-driven counterfactual reasoning by using a multi-objective optimization. We empirically apply our framework for finding a Pareto frontier of clustered optimal belief strengths that generalize across different contexts, demonstrating its efficacy on a toy dataset for credit decisions.

1. Introduction

With the rapid development of AI and ML and its pervasive use throughout all segments of society, there has been an increasing awareness of the necessity to assess and remediate the societal effects of these emerging technologies. Because

¹Google Deepmind ²The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA. Correspondence to: Michael S. Lee <ml5@andrew.cmu.edu>.

AI systems and their design will always encode some values or beliefs about the world, one major challenge to AI alignment is that some commonly accepted values have been shown to be fundamentally incompatible (Friedler et al., 2021).

For example, underlying value differences has led to an overwhelming variety of mathematical formulations of fairness metrics. But fairness impossibility theorems demonstrate that it is in fact mathematically impossible to simultaneously satisfy even the three common and intuitive definitions of fairness - demographic parity, equalized odds, and predictive rate parity (Miconi, 2017). Furthermore, individual fairness dictates that individuals who are similar should be treated similarly with respect to a specific outcome (Dwork et al., 2012). Yet group fairness states that demographic groups should, on the whole, receive similar decisions (Binns, 2020). And even if group fairness were the agreed upon standard, research has shown that even between groups, there are disparate notions of fairness and toxicity which are not universally generalizable. For example in the context of forum content moderation, each community implicitly subscribes to a different social contract, so toxicity for one community is not the same for another (Goyal et al., 2022). Although belief and value disagreement has increasingly become a critical issue in social alignment of new AI technology, it seems difficult to model, let alone resolve, these conflicts.

One way forward is spurred on by recent progress in large language models (LLMs), which has opened the door for AI researchers to develop general values via a set of language-based contracts (e.g. constitutions or rule sets) that are simultaneously-held without the need for direct (and perhaps conflicting) mathematical translation. Specifically, the technique aims to guide the training of an AI assistant through a list of rules and principles that guide the labeling of training data (Bai et al., 2022; Glaese et al., 2022; Thoppilan et al., 2022). Though rule-assisted methods have shown increased ability to provide harmless and helpful feedback, they each differ in their respective rule-sets due to subjective notions of good. Moreover, even if some set of constitutional tenets are generally accepted, each specific belief is held to varying degrees of strength among groups of individuals. And even in an individual, it is well known that belief uncertainty exists and is often not well calibrated, leading to the popular

overconfidence bias (Moore and Healy, 2008). This creates what we call the *meta-alignment problem*: Even if a general set of AI alignment values are identified, how should the model calibrate the strength of each value for beneficial societal impact?

Beneficial social impact is an inherent multi-faceted objective and has been captured by various dimensions of good, such as fairness, privacy, explainability, transparency, etc. Fairness efforts, such as those that equalize false positive rates for recidivism prediction (Dieterich et al., 2016), attempt to reduce bias and toxicity in algorithmic decisions or model outputs (Dwork et al., 2012). In a related vein, privacy efforts attempt to protect the sensitive attributes of an individual (Dwork, 2008) and other recent efforts like model cards (Mitchell et al., 2019) try to provide transparency into the model mechanisms as well as its benchmarked performance.

In our case, we focus on an overlooked yet crucial dimension of social benefit: cost-effective positive outcome generation, which is captured by the study of counterfactuals and recourses. Going back to the example of forum moderation, note that while high toxicity posts should be banned in an effort to maintain social order, it also can impede on the freedom of speech or expression. In light of this, it is thereby critical for the moderation system to minimize – or at least calibrate – the remediation effort it requires for an individual to post their honest opinions. Indeed, on another extreme, an AI moderation system that makes it impossible or extremely costly for a specific individual to express their opinions could be seen as a form of censorship and a violation of free speech (Massaro and Norton, 2015).

To address this, recent work has utilized counterfactual reasoning to consider not only different outcomes in response to different hypothetical input features (Wexler et al., 2019) or model choice (Bui et al., 2022) to minimize bias or overfitting (Gomez et al., 2021), but also in response to changes in features grounded in an individual in the form of recourses (Gomez et al., 2020). For example, linear models for credit determination has been analyzed for in the light of actionable recourse (Ustun et al., 2019), showing that given a negative outcome, the ability for an individual to induce counterfactual results often varies drastically based on model design. Therefore, recourse should be included as a critical factor to inform alignment efforts on producing positive societal impact.

Furthermore, counterfactuals have also long been studied in the social sciences as a tool in human decision making, for both learning from past outcomes (Roese, 1997), and weighing future outcomes to minimize future regret (Simonson, 1992). As we continue to build systems that ultimately strive to emulate human decision making, we argue for leveraging counterfactual reasoning to calibrate to human beliefs. And

as we deploy these systems in high impact domains such as finance (Murawski, 2019) and criminal justice (Hao, 2019), we argue for leveraging counterfactuals to adjust to foreseen and observed consequences in a principled manner both prior to and throughout deployment.

1.1. Our Contributions

In this paper, we provide a framework for tackling the meta-alignment problem of calibrating diverse human beliefs and values for maximizing beneficial societal impact, as defined from a counterfactual perspective. Building off of previous works that focus on outcome diversity from subjectivity among human populations, we argue that counterfactuals plays a critical role in calibrating and inducing belief diversity. Specifically, studies have shown the immense power that counterfactual reasoning (and anticipated regret) has to shape one’s beliefs *prior* to decision making, such as for avoiding risky behaviors like reckless driving. Furthermore, we note that the context in which these counterfactuals are analyzed is also critical for appropriately grounding the strength to which beliefs are upheld in decision making. For example, given the same classification task, human raters are shown to give more lenient labels when they are told that those labels are used for more severe penal punishments, suggesting that humans implicitly include the recourse cost in their calibration process.

In modeling belief diversity, we partition belief diversity into two categories: subjectivity and epistemic uncertainty. Subjectivity acknowledges that different individuals within a population maintain different sets of salient values to varying degrees. And epistemic uncertainty acknowledges that even an individual upholds these values to varying degrees across different contexts. The latter humbly acknowledges that an individual’s beliefs are generated by only a myopic view of the world, a view that is changing and shifting as that person constantly gains new insights and finds themselves in different circumstances (Schwartz, 2006). Specifically for epistemic uncertainty, we emphasize that counterfactuals help illuminate an important context in inducing belief diversity, as alignment values can change based on the task-dependent consequences.

Finally, in order to solve the two-fold problem of marrying alignment and calibration, we introduce a framework for understanding the belief update and calibration feedback loop, which allows for population and individual-level beliefs to not only govern decision-making policies but also to be updated based on the observed outcomes and recourses. This closed-loop cycle that leverages counterfactual reasoning to continuously calibrated the beliefs incorporated in our decision making algorithms, ultimately leading to desired outcomes and recourses (see Fig. 1). We leverage our Bayesian framework to not only induce outcome diversity

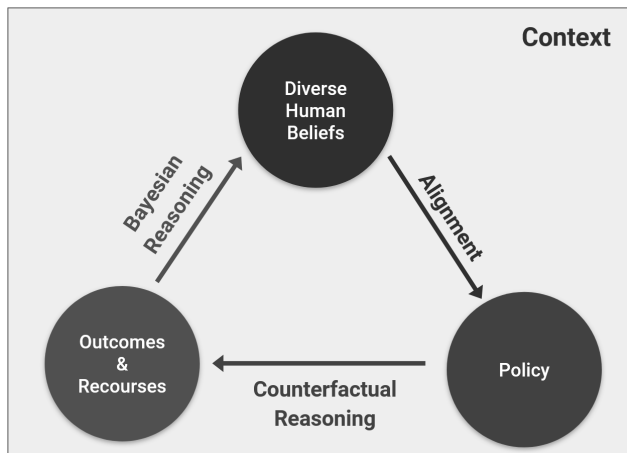



Figure 1. Belief calibration cycle. We propose to leverage counterfactual reasoning over the outcomes and recourses to calibrate towards beliefs that ensure outcomes and recourses that are beneficial toward decision subjects. Finally, we emphasize that this calibration occurs within a specific context that influences the beliefs, policy, outcomes and recourses and vice versa.

but also a distributional understanding of counterfactual recourses. We apply counterfactuals in multiple contexts to solve a multi-objective problem that attempts to model the multiple facets of societal impact. Unlike previous works in Bayesian modeling and alignment, we acknowledge that both the probabilistic accuracy of the predictions and the societal consequences of such predictions are critical for joint calibration.

We apply our contextual calibration approach on the `credit` dataset and introduce novel alignment contexts and measures of counterfactual recourse cost, such as the true negative and false negative recourse cost. For example, a loan business may wish to calibrate beliefs on the noise (σ) and sparsity (λ) model parameters by simultaneously 1) maximizing accuracy, 2) decreasing the recourse cost of loans that were falsely denied (false negative), while at the same time 3) increasing or thresholding on the recourse cost of loans that are justifiably denied (true negative). The last objective, while novel and perhaps counterintuitive in counterfactual cost calibration, is important for sensible alignment since access to free credit is not always desirable and costly consequences are critical for disincentivizing risky behavior. From our calibration experiments in this context, we find that only a few settings on belief priors on noise and sparsity lead to optimal outcomes in all three objectives (see Fig 2).

While there are multiple contexts and objectives, we find that not many of the suggested belief strengths are actually on the Pareto frontier. This suggests that although diversity in counterfactual contexts may lead to an explosion of valid

beliefs; in reality, there are only a small cluster of optimal calibrated beliefs. We empirically observe that diverse contextual information does in fact drive diversity in optimal belief strengths, often favoring more mild beliefs. Furthermore, our belief calibration framework reveals surprising trends in the task at hand, such as evidence that lower feature regularization does not always lead to lower recourse cost, or that increased credit leniency can help both predictive and social alignment.



σ	λ	FN Cost	TN Cost	Log-Prob
10	0.01	0.0088	0.0189	0.878
10	10	0.009	0.0189	0.866

Figure 2. Belief calibration on the `credit` dataset by minimizing average false negative recourse cost and negative log probability, via a scatter plot along with a table of Pareto optimal points. To disincentivize risky behavior, all beliefs corresponding to a higher risk of credit misuse (i.e. low true negative recourse cost) are removed, leaving only two optimal beliefs.

2. Subjectivity and Related Works

2.1. Early Efforts

Early AI approaches of symbolically modeling beliefs often divided the space into decisions that are either rational or irrational (Russell and Norvig, 2010), also referred to as cold and hot cognition (Tomkins et al., 1963). Abelson (1963) defined the problem as follows:

Is it possible to specify a realistic model for attitude change and resistance to change in sufficient process detail so that a computer could simulate it? [...] One might even speak of attitudinal problem-solving, wherein the individual is confronted with a challenge to his belief system and the “problem” he must solve is, “What am I to believe now?”

The major focus here is on cognitive processes, but Abelson proposes to explore some limited relations between cognition and affect. Specifically, within the context of attitudes and attitude changes, one might hope to develop a simulation model which would do for hot cognition what others have done for cold cognition. Implicit in this formulation is the idea that beliefs are fluid and are necessary to be updated to match the changing world.

These AI theorists posed that models of human intelligence were anchored to their belief systems, and that humans employed various techniques (like rationalization) to avoid changing the values of these systems, unlike the objective goals of winning a chess match (Abelson and Carroll, 1965). Systems like the Goldwater Machine (Wardrip-Fruin, 2009), POLITICS (Carbonell Jr, 1978), Terminal Time (Mateas et al., 2000), and RoleModel (Chen et al., 2010) were developed towards the pursuit of representing subjectivity (often due to affective and ideological differences) within humans.

2.2. Modern Efforts

As foundational as the Turing Test is among computational thought experiments (Turing, 1956), believability (or distinguishability) of human-like behavior has been a debatable measure of AI achievement (Searle, 2009). Recent advancements in LLMs, while trained to satisfy an objective score, can now believably represent diverse systems of beliefs and are a continuation of modeling the aforementioned “hot” aspects of our cognition, including the (1) subjectivity of each belief-system, and (2) the gravity that our myopia or epistemic miscalibration creates. Much work is focused on downstream remediation of more desirable model outputs, such as reward modeling from human feedback, or Reinforcement Learning from Human Feedback (RLHF) (Ouyang et al., 2022). Such approaches aim to bridge challenges in aligning to the nuances of human motivation. RLHF, for example, has shown promising results by directly providing feedback from human teachers that permits AI to learn with broader perspective and greater efficiency. However, such methods are nonetheless subject to the inherent biases of human behavior and focus mainly on predictive outcomes or increased personalization (Ziegler et al., 2019), rather than the consequences of such outcomes.

In light of this, as AI systems are technically more advanced, we posit that the tenuous dividing line between humans and AI, will hinge on analyzing the implicitly held values of an individual and the ethical consequences of such beliefs, as well as their ability to update their beliefs in a rational or irrational way. In some ways, the ability of an AI system to solve the meta-alignment problem of aligning the diversity of human values themselves is a modern day extension of Alan’s Turing test.

3. Belief Calibration via Contextual Counterfactuals

Counterfactual reasoning traditionally considers how altering a factual antecedent to an past event or decision may have lead to a subsequent change in outcome. In synthesizing the early research, Roese (1997) argues for a functional view of counterfactual thinking as a powerful tool for reasoning over past behavior to improve future outcomes. For example, counterfactuals are more likely to be generated in response to negative situations (Sanna and Turley, 1996), negative emotions (Roese, 1997), and repeated problems (Markman et al., 1993)).

3.1. Counterfactual Reasoning Calibrates Beliefs

As we design algorithms that increasingly impact peoples’ lives, we argue that the community should not only leverage counterfactuals to continuously learn from realized outcomes, but also that that they are a critical component in the re-calibration of our beliefs even prior to deployment. Studies have demonstrated the power that counterfactual reasoning (and anticipated regret) has to shape one’s beliefs *prior* to decision making. Parker et al. (1996) showed that out of videos that showing different possible outcomes of driving over the speed limit, it was specifically videos that focused on regret that most significantly changed their beliefs about unsafe driving. Similar results of a priori counterfactual reasoning affecting decision making can be found in consumers making safer choices in purchasing (Simonson, 1992) and unwillingness to trade identical lottery tickets for additional money, likely influenced by loss aversion toward regret that of trading away a winning ticket (Bar-Hillel and Neter, 1996).

As proposed by Roese, the power of counterfactual reasoning to promote changes in behavior lies especially in learning (post-hoc) or protecting against high consequence decisions (a priori). In 2016, Jeff Bezos popularized two types of decisions – Type 1 decisions that are consequential and effectively irreversible, in contrast to Type 2 decisions that are reversible – and argued that “great deliberation and consultation” are warranted only for the former (Bezos, 2016). We instead build on this framework by recognizing both consequence and reversibility both lie on spectrums and argue for on-going counterfactual reasoning in a calibration loop (see Fig. 1) over expected and outcomes and recourses (or reversibility) as we build algorithms that increasingly impact people’s lives.

3.2. Contextual Grounding Diversifies Beliefs

Though counterfactual reasoning is powerful and can help calibrate beliefs by closing the loop, the context in which it is done is critical for how strongly a belief should be up-

held. Context has long been acknowledged as being key for appropriately grounding decision making in affective computing (Pantic et al., 2006) and more recently in transparency and explainability (Ehsan et al., 2021), and can be broadly captured by the notion of the 5W’s that surround a decision, which are ‘who’ the involved parties are, ‘what’ the task is, and ‘where’, ‘when’, and ‘why’ the decision was made.

We illustrate the importance of grounding counterfactual reasoning in context through a recent study that highlighted the impact that perceived contextual consequences can have on changing human labeling beliefs. They found that humans will significantly differ in their labels of the same data depending on whether it is simply an exercise in factual description (e.g. whether a dog simply looked aggressive) or an exercise in normative judgment (e.g. whether a dog looked aggressive, and therefore violated an apartment’s policy) (Balagopalan et al., 2023). In discussing their results, the authors importantly note that participants may have attached “*a different cost to their judgments in the two conditions: Getting a decision wrong factually is just a matter of describing the world incorrectly. Getting it wrong normatively is a matter of potentially doing harm to another human.*” We posit that the ‘what’ of data labeling was shared between the two conditions with the identical set of features (e.g. aggression), but they differed in the ‘why’ (i.e. for description or judgment of a rule violation), and thus participants likely arrived at different counterfactuals of unequal consequence for the two conditions and adjusted their beliefs and labels accordingly.

Though different contexts may yield different counterfactuals that will influence human beliefs accordingly, shared contexts can conversely benefit from similar counterfactual reasoning that can yield a shared bias in beliefs. The aforementioned study by Balagopalan et al. (2023) also notes that participants in the judgment condition (or context) consistently showed “benefit of the doubt” in choosing not to flag close calls. Interestingly, this sentiment directly echoes the legal and judicial bias in the United States to usually favor the defendant, captured by the famous adage of “innocent until proven guilty” (Baradaran, 2011). Presumption of innocence is one of the long standing principles can be traced back to the religious and constitutional beliefs that founded America and aim to protect defendants from negative legal ramifications often with little recourse. Thus we may leverage operating in a shared context to bias beliefs in a shared direction that supports positive outcomes.

To conclude, we emphasize that people naturally use counterfactuals to assist them in calibrating their beliefs and decisions, especially in high stakes contexts with potentially negative outcomes. As we increasingly automate aspects of decision making that have significant ramification on

people’s livelihoods, we underscore the importance of also utilizing contextually-grounded counterfactual reasoning to calibrate the beliefs underlying our algorithms to help ensure beneficial outcomes for decision subjects.

4. Modeling Diversity of Beliefs

In this section, we give a Bayesian framework for modeling belief diversity mathematically and applying notions of social calibration for maximum impact. In Bayesian modeling, the idea of a belief is encapsulated by the prior distribution of a set of variables, as a proxy for the inductive and value biases of the model trainer. Generally, we extend this notion by also incorporating the choice of likelihood loss into the belief system structure, thereby reducing the beliefs and values of the trainer to be encapsulated more holistically by the posterior, or the trained model parameters. We note that policy beliefs, which guides the ultimate use of the model to make value-based judgments and decisions, are also integral but will not be mentioned here.

Specifically, let $(x_i, y_i) \in \mathcal{D}$ be datapoints and for model architecture parameterized by trainable parameters w , we see that $f_w(x)$ is typically trained by solving $w^* = \arg \min_w \mathcal{L}(f_w(x))$, where \mathcal{L} is some chosen objective with possibly many forms of regularization or penalties. The choice of objective, and also the training method, usually incorporates notions of the trainer’s belief to guide the parameters suitably, such as the preference for robustness in adversarial training (Gu and Rigazio, 2014), fairness-aware training via regularization on specific statistical parity based scores (Kamishima et al., 2011), or even the bias towards sparse solutions for ease of understanding and visualization (Rasmussen and Bro, 2012).

4.1. Subjectivity vs Epistemic Uncertainty

Different belief principles among large groups can be generally encapsulated by different loss functions $\mathcal{L}_1, \dots, \mathcal{L}_k$, leading to disparate trained parameters w_1, \dots, w_k . Oftentimes, these loss functions are a combination of fundamental beliefs, with its corresponding loss ℓ_i , as represented by the prior or the likelihood: $\mathcal{L} = \lambda_1 \ell_1 + \lambda_2 \ell_2 + \dots + \lambda_k \ell_k$, where ℓ_i are fundamental beliefs. For example, Bayesian LASSO (Park and Casella, 2008) trades off between the belief that prioritizes accuracy vs the belief for sparsity using some regularization parameter. Therefore, while subjectivity is the study of separate fundamental beliefs, we crucially note the strength to which each belief is expressed, or equivalently its certainty, is tuned as the study of epistemic uncertainty.

We apply Bayesian modeling to understand how uncertainty in beliefs translate over, via the model and its training process, to uncertainty in trained parameters, and ultimately to outcomes and recourses, where we focus on the for-

mer connection in this section. For simplicity, assume that $f_w(x) = \sum_i w_i \phi_i(x)$, where $\{\phi_i\}_{i=1}^k$ is some k dimensional feature map, with possibly $k = \infty$. Traditional Bayesian models can be viewed as transforming some prior beliefs about w to a posterior distribution over w , and is therefore the first step to understand the distribution over outcomes.

Generally, given some prior distribution $p(w)$ and by incorporating some noise beliefs about our observations, $y = f_w(x) + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma^2 \mathbb{I})$, we see that the posterior is derived by applying Bayes rule:

$$\begin{aligned} p(w|x_i, y_i) &\propto p(y|x_i, w)p(w) \\ &\propto \exp\left(-\frac{1}{2\sigma^2}\|y - f_w(x)\|^2\right)p(w) \\ &\propto \exp\left(-\frac{1}{2\sigma^2}\|y - \Phi(x)w\|^2\right)p(w) \end{aligned}$$

where $\Phi(x)$ is the matrix of $\phi_i(x)$. Therefore, we see that our likelihood is in fact a Gaussian with an inverse covariance matrix given by what is popularly known as the kernel matrix $K = \Phi(x)^\top \Phi(x)$, where for two input points $x, x' \in \mathbb{R}^d$, $K(x, x') = \sum_i \phi_i(x)\phi_i(x')$. Now, with the Gaussian prior $p(w)$, the resulting posterior is also Gaussian and is the basis of the popular Gaussian process regression (Rasmussen and Nickisch, 2010).

Even more generally, the noise model and the corresponding loss function may not correspond to a simple Gaussian distribution, as is often the case in classification; and furthermore, the prior may have a complicated expression that incorporates sophisticated beliefs. Therefore, the posterior is often not well analytically expressed, but Bayesian inference methods can extend to these regimes via sampling or approximation methods, such as MCMC and variational inference, usually according to the distribution $p(w|x_i, y_i) \propto \exp(-\mathcal{L}(f_w))p(w)$ (Dellaportas et al., 2002; Salimans et al., 2015).

4.2. Linear Regression

Since our model choice is targeted for maximal interpretability and explainability, we will focus on simplest notion of linear regression with $f_w(x) = w^\top x$ with $w \in \mathbb{R}^d$, in which case there is a direct correspondence from parameter to feature. Let us consider a per-feature prior on the parameters given by a general multi-variate Gaussian $p(w) = \mathcal{N}(m_p, \Sigma_p)$. Then applying Bayesian inference implies that the posterior is

$$p(w|x_i, y_i) \sim \mathcal{N}(w_{post}, A^{-1})$$

where $A = \sigma^{-2}XX^\top + \Sigma_p^{-1}$ and $w_{post} = \sigma^{-2}A^{-1}Xy$, where X, y are the stacked matrices of x_i, y_i .

Note that even this simple model, there are two classes of beliefs that are captured using the hyperparameters of posterior loss, parameterized by the 1) noise scale (σ) and 2) the regularization prior of the parameters (Σ_p). The noise scale measures the strength the belief that our data is inherently noisy, with $\sigma \rightarrow \infty$ implying that our data is completely random. The regularization prior Σ_p , when it is a diagonal matrix, represent a prior belief on how much features should be considered when making credit decisions, where $(\Sigma_p)_{ii} \rightarrow \infty$ implies that the feature i should not be used at all.

Typically, these hyperparameters are tuned using some sort of predictive calibration measures, such as the marginal log-likelihood. Automatic relevance determination (ARD) is a standard procedure for tuning Gaussian Process hyperparameters, with σ, Σ_p as the observation noise and length scale parameters (Rasmussen and Nickisch, 2010). However, while such measures give rise to accurate and well-calibrated models, they do not generally consider the downstream societal effects of these choices.

5. Calibration and Alignment

The two-fold problem of marrying alignment and calibration starts from the observation that 1) standard Bayesian calibration of hyperparameters is done without alignment to social beliefs and 2) mathematical alignment of models and policies to beliefs is often accomplished without calibration. Essentially, we either only calibrate to probabilistic measures of predictive success or to overgeneralized measures of ideal social outcomes, but not both. In this section, we give a preliminary solution to both issues by using a multi-objective approach for calibration and alignment, from the lens of counterfactual reasoning. Specifically, we consider the `credit` classification task that was used previously in the study on actionable recourse by linear regression (Ustun et al., 2019). We focus on experimental results, where we consider the calibration of a specific individual belief of feature-based weights via both outcome and recourse-based metrics, in multiple contexts.

The `credit` classification task aims to predict credit default using $d = 17$ features derived from their spending and payment patterns, education, credit history, age, and marital status. The dataset contains $n = 30000$ individuals and we transform the labels so that $y_i = -1$ if a person will default on an upcoming credit card payment and $y_i = 1$ otherwise. In our case, we can have diverse separate contexts: 1) recourse cost considers all features and 2) recourse cost considers only actionable features, i.e. those related to spending and payment patterns and education, or 3) recourse cost considers only unjustified negative credit decisions (false negative).

As mentioned before, we can represent our beliefs as hyperparameters on the prior parameterized by the 1) noise scale (σ) and 2) the regularization weights of the parameters (λ). The noise scale measures the strength the belief that incoming credit decisions are randomly made while the regularization weights λ represent a prior belief on how much features should be considered when making credit decisions. Generally, it is observed that as regularization increases, the number of individuals with recourse decreases as regularization reduces the number of actionable features.


5.1. Model Calibration

To calibrate both λ and σ , we formulate the search over different possible weights of $\lambda, \sigma \in [0.001, 0.01, 0.1, 1, 10]$ in an multi-objective optimization over with metrics that measure the consequences of such beliefs in the multiple contexts that we operate in. For the simplest setup, we consider λ as shared across all features and since we consider two metrics: log likelihood, a measure of predictive accuracy, and recourse cost, a measure of the minimal amount of effort needed to receive a positive classification with decent probability. We emphasize that our metrics are now inherently probabilistic, which highlights the power of our model to incorporate belief uncertainty to generate an outcome and recourse distribution.

Mathematically, our posterior distribution is given by $\mathcal{N}(w_{post}, A^{-1})$ which represents the diversity of beliefs arising from epistemic uncertainty. Then for an individual feature x , we note that $w^\top x \sim \mathcal{N}(w_{post}^\top x, x^\top A^{-1} x)$ is the predictive distribution. By using the sign as the decision threshold, the probability of receiving a positive outcome is given by the standard normal CDF of $\frac{w_{post}^\top x}{\sqrt{x^\top A^{-1} x}}$. Due to the low uncertainty estimates, we cap the log probability of wrong predictions at a minimum of -5 to prevent outliers from having exaggerated effects on aggregated statistics.

We fix our policy so that a positive credit decision is made if and only if $w_{post}^\top x \geq 0$, or equivalently if there is a positive outcome with at least 50% chance over all choices of the model weights. Thus, while each parameter provides a different recourse, we can define recourse to be the minimum norm of a correction needed to receive a positive classification with respect to our fixed policy. In this case, this is equivalent to solving the following optimization $cost = \min \|c\|$ s.t. $w_{post}^\top (x + c) \geq 0$. Note that this is 0 when $w_{post}^\top x \geq 0$.

Our empirical results (see Figure 3) show that there are many belief priors that produce suboptimal performance, as measured by our metrics. Concretely, we note that while we tested 25 discrete beliefs, there are only three that are Pareto optimal, specifically they are $\sigma = 10, \lambda = 0.1, 1, 10$. In fact, we emphasize that although all the Pareto optimal points



σ	λ	Recourse Cost	Log-Prob
10	0.1	0.0277	0.878
10	1	0.0279	0.865
10	10	0.028	0.866

Figure 3. Scatter plot of average recourse cost with negative log probability, along with a table of Pareto optimal points. Note that while many priors produce dominated points, there is a small Pareto frontier, which represents our set of optimal beliefs, and they have high confidence in this context.

share approximately the same negative log probability, the regularization value that produces the lowest recourse cost is $\lambda = 0.1$. From the Pareto frontier, it may seem that we agree with the conventional wisdom that as regularization increases, the average recourse cost increases since regularization reduces the number of actionable features. However, we note that the point with the highest recourse cost actually occurs when $\sigma = 10$ but $\lambda = 0.001$, which is perhaps surprising. These counter-intuitive conclusions highlight the impact and necessity of including counterfactual analysis in belief calibration.

5.2. Contextual Calibration

Appropriate definitions of recourse cost can vary significantly based on context. For example, we can use a feature-wise recourse cost to have lower costs for the actionable features, which are those related to spending and payment patterns and education. Specifically, let D be a diagonal matrix such that $D_{ii} = 1$ for actionable features but $D_{ii} = 100$ is large for non-actionable features. Then, the cost of recourse is solved by $cost(D) = \min \|Dc\|$ s.t. $w_{post}^\top (x + c) \geq 0$. We call this the actionable recourse cost. Accordingly, we also update our prior beliefs so that $\Sigma_p = D$, implying that we put higher emphasis on beliefs that non-actionable features should have near-zero weight and λ tunes the strength of this prior.




Figure 4. Scatter plot of average actionable recourse cost with negative log probability, where non-actionable features have substantially higher cost. Most of the Pareto frontier has not changed, however the strength parameters that gives rise to the lowest actionable recourse cost has a decreased value of σ, λ than before, justifying our argument that different contexts introduces more variance and uncertainty in Pareto optimal values of σ, λ .

In this context, we see that from Figure 4 that while most of the Pareto points are relatively unchanged, the belief that corresponds to the lowest actionable recourse cost now becomes $\sigma = 1, \lambda = 0.001$, as opposed to it being essentially $\sigma = 10$ from the previous analysis. In fact, in the previous context, $\sigma = 10$ holds for all Pareto optimal points, implying that there is a strong (and justified) belief that many of the credit decisions are noisy. However, by adding an additional context, this increased the epistemic uncertainty over the optimal values of σ as the much milder belief of $\sigma = 1$ is also now a candidate.

Another context that is particularly relevant for sensible alignment is to distinguish two forms of recourse cost by utilizing a novel split of the cost of remediating negative credit decisions into those that are justified (true negative) and those that are unjustified (false negative). While decreasing actionable recourse has been extensively studied, especially when considering unjustified credit denials, we posit that it is perhaps equally important to uphold sufficiently costly consequences for justified credit denials in order to achieve societal alignment that disincentivizes risky behavior. Specifically, we define the false negative cost as equal to the recourse cost when the label is in fact 1 and the

cost is 0 otherwise. Analogously, we define the true negative cost and our main objective is to minimize the false negative cost, while maintaining an adjusted minimum on the true negative cost, a threshold that we empirically set at 0.0188. Our calibration results are shown in Figure 2 and we find that only two reasonable values of σ, λ are proposed.

Overall, we observe that introducing new, yet related, contexts do not create substantially wild variation in beliefs, but simply proposes a mild re-calibration, or a strict subset of common belief settings that generalize across contexts.. Specifically, note that some of our previously optimal beliefs are still on the altered Pareto frontier under diverse contexts (note also that when adding a new context as an additional objective, it will not remove points on the Pareto front). This implies that our beliefs on relative prior strengths generalize across different contexts, but introducing differing costs can justify adding another belief prior into our set of non-dominated beliefs. However, we note that belief calibration can differ significantly given new counterfactual information.

5.3. Policy Calibration

Model design and calibration is only one of multiple processes that can be aligned in the credit outcome decision. In this section, by introducing subjectivity, we add a new hyperparameter β , that captures beliefs in policy design, which takes a model output to outcomes and recourses. Just as in Bayesian optimization, we can exploit our model’s uncertainty predictions to encode the belief of giving individuals the “benefit of the doubt” by giving a positive outcome when $\mu(x) + \beta\sigma(x) \geq 0$, where μ, σ are the predicted mean and standard deviation. Note that previously our policy is setting $\beta = 0$ and therefore, our more lenient policy is equivalent to allowing easier credit when there is higher uncertainty or doubt in the model.

Mathematically, we now define recourse to be a slightly complicated quadratic optimization of the form $cost(\beta) = \min \|c\|$ s.t. $w_{post}^\top(x + c) + \beta(x + c)^\top A^{-1}(x + c) \geq 0$. This requires an eigenvector computation and is significantly more computationally expensive; therefore we use a subset of the dataset, around 1000 individuals, to perform our experiments. Furthermore, we note that the log-likelihood calculations now include β term, to reflect the greater leniency to generating a positive outcome. We let $\beta \in [0.1, 1, 10]$, where a larger β represents easier credit.

From our results, we find the surprising fact that giving more lenient credit decisions can actually help both accuracy while reducing policy-based recourse cost (see Fig 5). Specifically, the Pareto optimal beliefs are calibrated to have high values of σ, β , implying that higher uncertainty in credit modeling and using those uncertainty measures in giving more lenient credit decisions improve our predictive

calibration, as the most accurate model has a $\beta = 1$ and $\sigma = 10$.




Figure 5. Scatter plot of average policy-based recourse cost with negative log probability, along with a table of Pareto optimal points. Note that all the Pareto frontiers do not use the small values of σ, β , meaning that calibrating our beliefs to recognize that credit decisions are inherently noisy and adding some leniency into the final decision policy produces more accurate probabilistic predictions while also promoting better societal outcomes.

6. Conclusion

As we continue to build decision systems modeled after human beliefs in high stakes contexts, it is critical to acknowledge the diversity of human beliefs, rising from both population-level subjectivity but also individual-level epistemic uncertainty, and the difficulty that such diversity presents, which we term the meta-alignment problem. To tackle the challenge of calibrating beliefs to different contexts, we proposed to leverage counterfactual reasoning to not only consider measures of predictive success but also social benefit in the form of outcomes and recourses in a “belief calibration cycle.” We presented experiments on a credit decision dataset, leveraging Bayesian modeling to analyze how uncertainty in beliefs on the hyperparameter values leads to a distribution of outcomes and recourses, which in turn informed the optimal choices in beliefs for obtain the best outcomes and recourses. By doing so, we find that belief calibration leads to surprising conclusions, such as the case for increased leniency in credit applications, which highlights the importance of utilizing counterfactual

analysis in belief calibration.

References

- R. P. Abelson. Computer simulation of “hot cognition”. *Computer simulation of personality*, pages 277–298, 1963.
- R. P. Abelson and J. D. Carroll. Computer simulation of individual belief systems. *American behavioral scientist*, 8(9):24–30, 1965.
- Y. Bai, S. Kadavath, S. Kundu, A. Askell, J. Kernion, A. Jones, A. Chen, A. Goldie, A. Mirhoseini, C. McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022.
- A. Balagopalan, D. Madras, D. H. Yang, D. Hadfield-Menell, G. K. Hadfield, and M. Ghassemi. Judging facts, judging norms: Training machine learning models to judge humans requires a modified approach to labeling data. *Science Advances*, 9(19):eabq0701, 2023.
- M. Bar-Hillel and E. Neter. Why are people reluctant to exchange lottery tickets? *Journal of Personality and Social Psychology*, 70(1):17, 1996.
- S. Baradaran. Restoring the presumption of innocence. *Ohio St. LJ*, 72:723, 2011.
- J. Bezos. Letter to shareholders, 2016. URL <https://www.sec.gov/Archives/edgar/data/1018724/000119312516530910/d168744dex991.htm>.
- R. Binns. On the apparent conflict between individual and group fairness. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 514–524, 2020.
- N. Bui, D. Nguyen, and V. A. Nguyen. Counterfactual plans under distributional ambiguity. In *International Conference on Learning Representations*, 2022.
- J. G. Carbonell Jr. Politics: Automated ideological reasoning. *Cognitive Science*, 2(1):27–51, 1978.
- S. Chen, A. M. Smith, A. Jhala, N. Wardrip-Fruin, and M. Mateas. Rolemodel: towards a formal model of dramatic roles for story generation. In *Proceedings of the Intelligent Narrative Technologies III Workshop*, pages 1–8, 2010.
- P. Dellaportas, J. J. Forster, and I. Ntzoufras. On bayesian model and variable selection using mcmc. *Statistics and computing*, 12(1):27–36, 2002.

- W. Dieterich, C. Mendoza, and T. Brennan. Compas risk scales: Demonstrating accuracy equity and predictive parity. *Northpointe Inc*, 7(4), 2016.
- C. Dwork. Differential privacy: A survey of results. In *Theory and Applications of Models of Computation: 5th International Conference, TAMC 2008, Xi'an, China, April 25-29, 2008. Proceedings 5*, pages 1–19. Springer, 2008.
- C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226, 2012.
- U. Ehsan, Q. V. Liao, M. Muller, M. O. Riedl, and J. D. Weisz. Expanding explainability: Towards social transparency in ai systems. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–19, 2021.
- S. A. Friedler, C. Scheidegger, and S. Venkatasubramanian. The (im) possibility of fairness: Different value systems require different mechanisms for fair decision making. *Communications of the ACM*, 64(4):136–143, 2021.
- A. Glaese, N. McAleese, M. Trebacz, J. Aslanides, V. Firoiu, T. Ewalds, M. Rauh, L. Weidinger, M. Chadwick, P. Thacker, et al. Improving alignment of dialogue agents via targeted human judgements. *arXiv preprint arXiv:2209.14375*, 2022.
- O. Gomez, S. Holter, J. Yuan, and E. Bertini. Vice: Visual counterfactual explanations for machine learning models. In *Proceedings of the 25th International Conference on Intelligent User Interfaces*, pages 531–535, 2020.
- O. Gomez, S. Holter, J. Yuan, and E. Bertini. Advice: Aggregated visual counterfactual explanations for machine learning model validation. In *2021 IEEE Visualization Conference (VIS)*, pages 31–35. IEEE, 2021.
- N. Goyal, I. Kivlichan, R. Rosen, and L. Vasserman. Is your toxicity my toxicity? exploring the impact of rater identity on toxicity annotation. *arXiv preprint arXiv:2205.00501*, 2022.
- S. Gu and L. Rigazio. Towards deep neural network architectures robust to adversarial examples. *arXiv preprint arXiv:1412.5068*, 2014.
- K. Hao. Ai is sending people to jail—and getting it wrong, 2019. URL <https://www.technologyreview.com/2019/01/21/137783/algorithms-criminal-justice-ai/>.
- T. Kamishima, S. Akaho, and J. Sakuma. Fairness-aware learning through regularization approach. In *2011 IEEE 11th International Conference on Data Mining Workshops*, pages 643–650. IEEE, 2011.
- K. D. Markman, I. Gavanski, S. J. Sherman, and M. N. McMullen. The mental simulation of better and worse possible worlds. *Journal of experimental social psychology*, 29(1):87–109, 1993.
- T. M. Massaro and H. Norton. Siri-ously? free speech rights and artificial intelligence. *Nw. UL Rev.*, 110:1169, 2015.
- M. Mateas, P. Vanouse, and S. Domike. Generation of ideologically-biased historical documentaries. 2000.
- T. Miconi. The impossibility of “fairness”: a generalized impossibility result for decisions. *arXiv preprint arXiv:1707.01195*, 2017.
- M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji, and T. Gebru. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 220–229, 2019.
- D. A. Moore and P. J. Healy. The trouble with overconfidence. *Psychological review*, 115(2):502, 2008.
- J. Murawski. Mortgage providers look to ai to process home loans faster, 2019. URL <https://www.wsj.com/articles/mortgage-providers-look-to-ai-to-process-home-loans-faster-11552899212>.
- L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- M. Pantic, A. Pentland, A. Nijholt, and T. Huang. Human computing and machine understanding of human behavior: A survey. In *Proceedings of the 8th international conference on Multimodal interfaces*, pages 239–248, 2006.
- T. Park and G. Casella. The bayesian lasso. *Journal of the American Statistical Association*, 103(482):681–686, 2008.
- D. Parker, S. G. Stradling, and A. S. Manstead. Modifying beliefs and attitudes to exceeding the speed limit: an intervention study based on the theory of planned behavior 1. *Journal of Applied Social Psychology*, 26(1):1–19, 1996.
- C. E. Rasmussen and H. Nickisch. Gaussian processes for machine learning (gpml) toolbox. *Journal of machine learning research*, 11(Nov):3011–3015, 2010.

- M. A. Rasmussen and R. Bro. A tutorial on the lasso approach to sparse modeling. *Chemometrics and Intelligent Laboratory Systems*, 119:21–31, 2012.
- N. J. Roese. Counterfactual thinking. *Psychological bulletin*, 121(1):133, 1997.
- S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 3 edition, 2010.
- T. Salimans, D. Kingma, and M. Welling. Markov chain monte carlo and variational inference: Bridging the gap. In *International conference on machine learning*, pages 1218–1226. PMLR, 2015.
- L. J. Sanna and K. J. Turley. Antecedents to spontaneous counterfactual thinking: Effects of expectancy violation and outcome valence. *Personality and Social Psychology Bulletin*, 22(9):906–919, 1996.
- S. H. Schwartz. Basic human values: An overview. 2006.
- J. Searle. Chinese room argument. *Scholarpedia*, 4(8):3100, 2009.
- I. Simonson. The influence of anticipating regret and responsibility on purchase decisions. *Journal of consumer research*, 19(1):105–118, 1992.
- R. Thoppilan, D. De Freitas, J. Hall, N. Shazeer, A. Kulshreshtha, H.-T. Cheng, A. Jin, T. Bos, L. Baker, Y. Du, et al. Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*, 2022.
- S. Tomkins, S. Messick, E. T. Service, and P. University. *Computer Simulation of Personality: Frontier of Psychological Theory*. Wiley, 1963. ISBN 9780598635334. URL <https://books.google.com/books?id=M.ZGAAAAMAAJ>.
- A. M. Turing. Can a machine think. *The world of mathematics*, 4:2099–2123, 1956.
- B. Ustun, A. Spangher, and Y. Liu. Actionable recourse in linear classification. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 10–19, 2019.
- N. Wardrip-Fruin. *Expressive Processing: Digital fictions, computer games, and software studies*. MIT press, 2009.
- J. Wexler, M. Pushkarna, T. Bolukbasi, M. Wattenberg, F. Viégas, and J. Wilson. The what-if tool: Interactive probing of machine learning models. *IEEE transactions on visualization and computer graphics*, 26(1):56–65, 2019.
- D. M. Ziegler, N. Stiennon, J. Wu, T. B. Brown, A. Radford, D. Amodei, P. Christiano, and G. Irving. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*, 2019.