

EmoDebt: Bayesian-Optimized Emotional Intelligence for Strategic Agent-to-Agent Debt Recovery

Yunbo Long

University of Cambridge
Cambridge, United Kingdom
yl1892@cam.ac.uk

Liming Xu

University of Cambridge
Cambridge, United Kingdom
lx249@cam.ac.uk

Yuhan Liu

University of Toronto
Toronto, Canada
yuhanlydia.liu@alumni.utoronto.ca

Alexandra Brintrup

University of Cambridge
Cambridge, United Kingdom
ab702@cam.ac.uk

ABSTRACT

The rise of autonomous LLM agents has enabled strategic agent-to-agent interactions, yet a critical vulnerability persists: in high-stakes, emotion-sensitive domains like debt collection, LLM agents pre-trained on human dialogue are vulnerable to exploitation by adversarial counterparts who simulate negative emotions to derail negotiations. To fill this gap, we first contribute a novel dataset of simulated debt recovery scenarios and a multi-agent simulation framework. Within this framework, we introduce EmoDebt, an LLM agent architected for robust performance. Its core innovation is a Bayesian-optimized emotional intelligence engine that reframes a model’s ability to express emotion in negotiation as a sequential decision-making problem. Through online learning, this engine continuously tunes EmoDebt’s emotional transition policies, discovering optimal counter-strategies against specific debtor tactics. Extensive experiments on our proposed benchmark demonstrate that EmoDebt achieves significant strategic robustness, substantially outperforming non-adaptive and emotion-agnostic baselines across key performance metrics, including success rate and operational efficiency. By introducing both a critical benchmark and a robustly adaptive agent, this work establishes a new foundation for deploying strategically LLM agents in adversarial debt interactions. The code is available at <https://github.com/Yunbo-max/EmoDebt>.

KEYWORDS

LLM Agents; Multi-turn Negotiation; Bayesian Optimization

ACM Reference Format:

Yunbo Long, Yuhan Liu, Liming Xu, and Alexandra Brintrup. 2026. EmoDebt: Bayesian-Optimized Emotional Intelligence for Strategic Agent-to-Agent Debt Recovery. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 9 pages. <https://doi.org/10.65109/OCSI2038>

1 INTRODUCTION

Credit finance underpins modern economies, with effective debt collection being a critical yet complex component [18]. Traditional

models recognize debt recovery not merely as a financial transaction but as a profoundly sensitive interpersonal process. These dialogues are emotionally charged by nature, as they involve financial distress, personal circumstances, and reputational consequences, requiring a careful balance of empathy, trust, and strategic pressure [3, 19]. Research in affective computing has demonstrated that emotional intelligence can foster trust and cooperation in financial negotiations [4, 30], while large financial institutions have recognized that embedding empathy into negotiation processes is essential for sustaining client relationships and achieving long-term success [7, 15]. The emergence of Large Language Models (LLMs) offered a promising path toward automating these emotionally complex interactions through AI-powered systems [7, 24]. However, a significant limitation persists: these systems are largely confined to static, reactive dialogues and lack the sophisticated, adaptive reasoning required for the dynamic strategic play of a genuine negotiation [23]. This gap is exacerbated by the stringent privacy of financial data, which has historically prevented the creation of public, large-scale debt recovery datasets needed to train and benchmark such systems [6]. In addition, a shift is now underway, moving beyond human-in-the-loop systems toward a new paradigm of direct agent-to-agent interaction [14, 22]. This emerging multi-agent ecosystem enables complex, high-volume transactions in fields like e-commerce and decentralized finance (DeFi) autonomously [28]. Debt recovery is a natural fit for this automation, as it involves high-frequency, protocol-driven interactions that are costly to scale with human agents. Besides, it creates an ideal testbed for developing and benchmarking strategic reasoning in LLM agents against adversarial counterparts. However, this transition to an agent-to-agent paradigm introduces a critical and previously unaddressed vulnerability. The core strength of LLMs—their training on vast corpora of human dialogue, making them proficient in parsing and generating emotionally-laden language—becomes a profound liability in an adversarial setting [9, 21]. A strategic debtor agent can weaponize this capability, deploying simulated emotional states like strategic anger, fabricated distress, or tactical indignation. These are not genuine expressions but calculated ploys designed to exploit the creditor agent’s empathetic programming, derailing negotiations, extracting unjustified concessions, and prolonging recovery cycles. This creates a pressing need for strategically robust agents that can operate effectively, where emotional intelligence is no longer about empathy but about defense against strategic manipulation.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/OCSI2038>

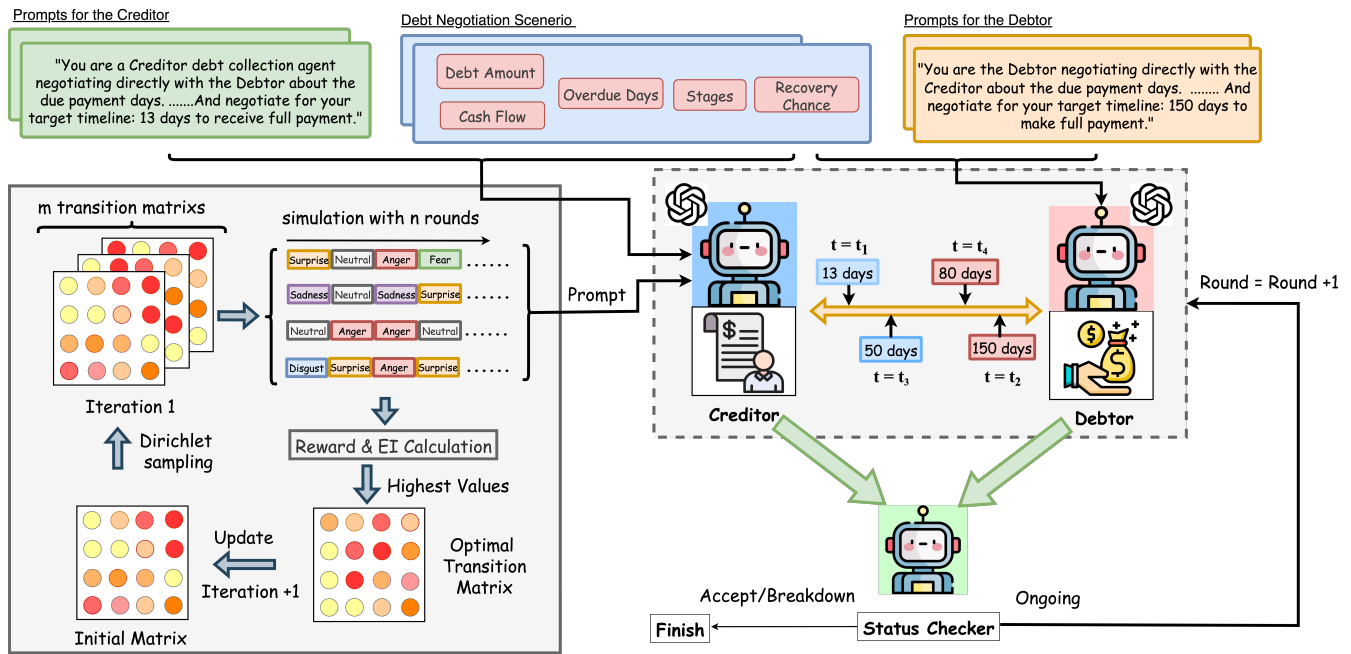


Figure 1: Illustration of the Pipeline of EmoDebt.

To fill this gap, our first contribution is the creation of a novel debt recovery dataset and a multi-agent simulation framework. This platform overcomes the scarcity of public financial negotiation data by providing realistic scenarios with diverse debtor profiles, built on the LangGraph architecture to enable controlled interactions between autonomous creditor and debtor agents. Within this framework, we introduce EmoDebt, a novel LLM agent designed for strategic robustness in debt negotiation. EmoDebt’s core innovation is a Bayesian-optimized emotional intelligence engine, which reframes emotional intelligence as a sequential decision-making problem. We formalize emotional strategy through a 7×7 transition probability matrix across seven emotional states, initialized with psychologically informed priors. A Gaussian Process-based Bayesian optimizer then treats negotiation outcomes as a black-box function, continuously learning optimal emotional transition policies. Through an online reinforcement mechanism that rewards successful agreements, favorable payment terms, and negotiation efficiency, EmoDebt dynamically discovers high-performing emotional counter-strategies.

We validate our approach through comprehensive experiments across multiple LLM configurations. The results demonstrate that EmoDebt achieves a near-perfect 99.7% success rate in optimal settings and an average improvement of +46.2% in success rate across all model pairings. Furthermore, it can reduce collection timelines and negotiation duration by 86.5% and 67.5%, respectively, compared to a pure agent-to-agent setting without emotional guidance for the creditor. This establishes a novel system for emotionally adaptive negotiation. The main contributions of this work are three-fold: (1) the development of a novel *debt recovery dataset and multi-agent simulation framework* for controlled, repeatable negotiation

experiments; (2) the introduction of **EmoDebt**, an LLM agent that uses a *Bayesian-optimized emotional intelligence engine* to reframe emotional expression as a sequential decision-making problem for strategic robustness; and (3) empirical validation demonstrating that EmoDebt achieves a near-perfect success rate while significantly reducing collection time and negotiation turns compared to non-emotional baselines.

2 RELATED WORK

2.1 Emotional Intelligence in Debt Collection

Researchers have emphasized the critical role of emotional intelligence in debt collection [11, 12]. For instance, Bachman et al. [2] conducted a comparative study on debt collectors’ performance using the Emotional Quotient Inventory (EQ-i). Their findings revealed that account officers with higher emotional intelligence tend to achieve better job performance. Similarly, Liao et al. [10] explored the effectiveness of debt collection strategies by analyzing voice and text data from debt collection phone calls. Their research demonstrated that strategies evoking happiness and fear are effective in reducing repayment time, whereas those that diminish happiness or provoke anger hinder repayment. These findings highlight the importance of considering the emotional impact on debtors when performing debt collection.

2.2 LLMs in Automated Negotiation

While LLMs show promise in automating debt negotiation, few studies have examined their emotional strategies in this context. Schneider et al. [23] applied LLMs to price negotiations with humans but overlooked the role of emotional dynamics. Similarly, Wang et al. [27] found that LLMs tend to over-concede compared

to human negotiators and proposed a Multi-Agent approach to improve decision rationality, yet they did not account for the function of emotions in negotiation. Typically, LLM-based agents mimic empathy by recognizing patterns in their training data rather than employing strategic emotional reasoning. Without genuine affective understanding, they struggle to adjust their tone and negotiation strategy based on a debtor’s emotional state. For instances, if a debtor gets angry, the agent may escalate tension. And if the debtor sound desperate, the agent might concede unfairly [1]. Lacking emotional intelligence, agents may fail to effectively navigate real-world negotiations.

2.3 Agent-to-Agent Negotiation Systems

The field of automated negotiation has seen significant development in agent-to-agent interaction frameworks. Foundational work by [31] established protocols for multi-issue negotiation between autonomous agents. More recently, [16] demonstrated how deep reinforcement learning can be applied to create agents that develop sophisticated bargaining strategies through self-play. These agent-based systems provide valuable testbeds for evaluating negotiation algorithms but typically operate in emotion-free environments[17], limiting their applicability to human-facing scenarios like debt collection where emotional factors are crucial.

2.4 Strategy Learning in Multi-Agent Systems

Research on strategy learning in multi-agent systems offers relevant methodologies for improving negotiation agents[13]. Online learning approaches, such as those explored by [8], allow agents to adapt their strategies based on real-time interaction outcomes. [20] applied deep reinforcement learning to develop agents that can learn effective negotiation policies through repeated interactions. These computational frameworks demonstrate the potential for creating adaptive negotiation agents[4], but have not been sufficiently integrated with emotional intelligence components necessary for finance applications like debt collection.

3 EMODEBT

We formulate the debt collection negotiation as a sequential decision-making problem between two autonomous LLM agents: a creditor agent C and a debtor agent D . As shown in 1, the negotiation proceeds in discrete rounds $t = 1, 2, \dots, n$, where n is the maximum dialog length. It includes (1) Debt Negotiation Setup: We configure creditor and debtor agents using state-of-the-art LLMs, providing debt details (Overdue days, Debt amount, Cash Flow and so on). (2) EmoDebt Optimization: We evolve creditor emotion policies through Bayesian optimization. Each iteration tests emotional transition matrices, evaluates them via multi-turn debt negotiations using reward function based on collection days and success rate, then calculates Expected Importance (EI) to optimize emotional transitions. This process iterates until convergence, yielding optimal creditor emotion strategies.

3.1 Problem Formulation

Let \mathcal{S} be the state space representing negotiation progress (e.g., offer, acceptance, breakdown), and $\mathcal{A}_C, \mathcal{A}_D$ be the action spaces for creditor and debtor respectively, comprising possible negotiation

messages and emotional responses. The negotiation evolves as a Markov Decision Process $s_{t+1} = f(s_t, a_t^C, a_t^D)$, which defines the state transition dynamics, where the next negotiation state s_{t+1} depends on the current state s_t and the joint actions of both agents (a_t^C, a_t^D) . The function f encapsulates the complex interaction dynamics between the LLM-based agents.

3.2 Emotional State Modeling

We model the creditor’s emotional strategy using a finite set of emotional states $\mathcal{E} = \{e_1, e_2, \dots, e_7\}$, including *happy, surprising, angry, sad, disgust, fear, and neutral*. This comprehensive set captures the full spectrum of strategic emotional responses relevant to debt collection scenarios. The emotional transitions are governed by a stochastic policy represented as a transition probability matrix $\mathbf{P} \in \mathbb{R}^{7 \times 7}$:

$$\mathbf{P}_{ij} = \mathbb{P}(e_{t+1} = e_j \mid e_t = e_i), \quad (1)$$

which defines the core emotional transition model, where \mathbf{P}_{ij} represents the probability of transitioning from emotional state e_i to emotional state e_j in the next negotiation round. This stochastic formulation allows for flexible emotional adaptation while maintaining psychological plausibility in emotional sequencing. The transition matrix must satisfy the probability constraints as $\sum_{j=1}^7 \mathbf{P}_{ij} = 1$ for all $i \in \{1, \dots, 7\}$, where the normalization constraint ensures that from any emotional state e_i , the probabilities of transitioning to all possible next states sum to 1. The emotional transition matrix is initialized using psychologically-grounded priors (Table 2) based on [25, 26]. Specifically, $\mathbf{P}_{ij}^{(0)} = \pi_{ij}$, where each prior value π_{ij} represents established transitional probabilities between emotions.

3.3 Bayesian Optimization Framework

We treat the negotiation outcome as a black-box function $g : \mathbb{R}^{49} \rightarrow \mathbb{R}$ that maps the flattened transition matrix $\mathbf{p} = \text{vec}(\mathbf{P})$ to a reward:

$$r = g(\mathbf{p}), \quad (2)$$

which frames our optimization problem, where the 49-dimensional vector \mathbf{p} (flattened 7×7 transition matrix) serves as input to an unknown reward function g . This black-box formulation is appropriate because the relationship between emotional strategies and negotiation outcomes is complex and non-linear. The reward function combines debt recovery efficiency and negotiation speed:

$$r(\mathbf{p}) = \begin{cases} -\alpha \cdot \log(n_{\text{rounds}}) / d_{\text{extended}} & \text{if successful negotiation} \\ -d_{\text{max}} & \text{if negotiation fails} \end{cases}, \quad (3)$$

which defines our reward function where successful negotiations are penalized by the product of final extended collection days d_{ex} and logarithmic round count $\log(n_{\text{rounds}})$, scaled by α . Failed negotiations receive a fixed penalty of maximum debt days d_{max} , maintaining negative scaling while emphasizing both timeline length and negotiation efficiency. Besides, we model the unknown function g using a Gaussian Process (GP) due to its sample efficiency and uncertainty quantification capabilities:

$$g(\mathbf{p}) \sim \mathcal{GP}(\mu(\mathbf{p}), k(\mathbf{p}, \mathbf{p}')), \quad (4)$$

which places a Gaussian Process prior over the reward function, where $\mu(\mathbf{p})$ is the mean function (typically set to zero after normalization) and $k(\mathbf{p}, \mathbf{p}')$ is the covariance kernel that encodes our

assumptions about function smoothness and correlation structure. We employ the Matérn kernel for its flexibility in modeling various smoothness regimes:

$$k(\mathbf{p}, \mathbf{p}') = \sigma_f^2 \left(1 + \frac{\sqrt{3}\|\mathbf{p} - \mathbf{p}'\|}{\ell} \right) \exp \left(-\frac{\sqrt{3}\|\mathbf{p} - \mathbf{p}'\|}{\ell} \right), \quad (5)$$

which specifies the Matérn 3/2 kernel, where σ_f^2 controls the function variance, ℓ is the length-scale parameter determining the correlation distance, and $\|\mathbf{p} - \mathbf{p}'\|$ is the Euclidean distance between emotional strategy vectors. This kernel is particularly suitable for emotional dynamics as it assumes only once-differentiable functions, matching realistic emotional transition patterns.

3.4 Online Learning

At each iteration k , given historical observations $\mathcal{D}_k = \{(\mathbf{p}_i, r_i)\}_{i=1}^k$ where $\mathbf{p}_i = \text{vec}(\mathbf{P}_i)$ represents the flattened emotional transition matrix and r_i is the corresponding reward, we employ Expected Improvement (EI) as the acquisition function to balance exploration and exploitation:

$$\text{EI}(\mathbf{p}) = \mathbb{E}[\max(0, g(\mathbf{p}) - g(\mathbf{p}^+) - \xi)], \quad (6)$$

where \mathbf{p} represents candidate emotional strategy (flattened 49-dimensional vector). And $g(\mathbf{p})$ means Gaussian Process prediction of reward for strategy \mathbf{p} . In addition, $g(\mathbf{p}^+)$ shows best reward value observed in history \mathcal{D}_k . And item $\xi = 0.01$ is the exploration parameter that controls risk. The EI criterion quantifies the expected improvement over the current best strategy, where higher values indicate more promising regions of the emotional strategy space to explore. Candidate emotional transition matrices are generated via Dirichlet perturbations to ensure valid probability distributions while maintaining psychological plausibility:

$$\mathbf{P}_{\text{candidate}}^{(i)} \sim \text{Dirichlet}(\alpha \cdot \mathbf{P}_{\text{current}}^{(\text{row})} + \epsilon), \quad (7)$$

where $\mathbf{P}_{\text{current}}^{(\text{row})}$ represents current emotional transition matrix row probabilities, $\alpha = 10.0$ acts as the concentration parameter controlling perturbation magnitude, with higher values producing candidates that closely resemble the current strategy (exploitation) and lower values enabling more diverse candidate generation (exploration). The parameter $\epsilon = 0.1$ serves as a smoothing constant for numerical stability, preventing zero probabilities in Dirichlet sampling and ensuring all emotional transitions remain possible.

The Bayesian optimization update selects the most promising emotional strategy through:

$$\mathbf{P}_{k+1} = \arg \max_{\mathbf{P} \in C_k} \text{EI}(\text{vec}(\mathbf{P})), \quad (8)$$

where C_k denotes the set of candidate matrices generated via Equation (10), $\text{EI}(\text{vec}(\mathbf{P}))$ represents the Expected Improvement for candidate \mathbf{P} , and \mathbf{P}_{k+1} indicates the selected emotional strategy for the next iteration. This systematic approach enables efficient exploration of the high-dimensional emotional strategy space (49 dimensions) while strategically leveraging historical negotiation outcomes to focus on promising regions. The Dirichlet perturbations ensure that all candidate matrices maintain valid probability distributions ($\sum_j P_{ij} = 1$), while the EI acquisition function directs the search toward strategies that either show high predicted reward (exploitation) or high uncertainty (exploration).

3.5 Theoretical Guarantees

Under the assumption that the reward function g is Lipschitz continuous [5] and the emotional transition space is compact, our Bayesian optimization approach achieves asymptotic convergence:

$$\lim_{k \rightarrow \infty} \mathbb{P}(g(\mathbf{p}_k) \geq g(\mathbf{p}^*) - \epsilon) = 1 \quad (9)$$

, which provides the theoretical guarantee that as the number of iterations increases, the probability that our discovered emotional strategy \mathbf{p}_k achieves reward within ϵ of the global optimum \mathbf{p}^* approaches 1. This ensures that with sufficient negotiation experience, EmoDebt will converge to near-optimal emotional response patterns. Besides, we monitor exploration diversity using emotional transition matrix entropy:

$$H(\mathbf{P}) = -\frac{1}{7} \sum_{i=1}^7 \sum_{j=1}^7 P_{ij} \log P_{ij}, \quad (10)$$

which defines the normalized entropy [29] of the emotional transition matrix, which serves as a diagnostic measure for our learning process. High entropy indicates diverse emotional exploration, while low entropy suggests convergence to specific emotional patterns. This metric helps balance exploration of new emotional strategies against exploitation of known effective ones.

3.6 Multi-Agent Simulation Framework

Our multi-agent simulation framework comprises three specialized LLM agents that interact in a controlled debt collection environment as shown in Algorithm 1: the Creditor Agent (\mathcal{M}_C) implements the EmoDebt emotional intelligence engine using Bayesian-optimized transition matrices to generate emotionally-aware responses; the Debtor Agent (\mathcal{M}_D) simulates realistic debtor behavior with configurable emotional strategies including angry, sad, fearful, and manipulative profiles; and the Examiner Agent (\mathcal{M}_E) monitors negotiation progress, detects terminal states (accept/breakdown), and computes performance metrics for reward evaluation. Through iterative dialogue rounds where emotional states evolve stochastically, this framework enables large-scale testing of emotional strategies across diverse scenarios, providing the empirical foundation for Bayesian optimization while ensuring reproducible evaluation of emotional intelligence in autonomous debt collection negotiations.

4 EXPERIMENTS

4.1 Debt Dataset.

For our experiments, we created the Credit Recovery Assessment Dataset (CRAD), comprising 100 synthetic credit delinquency cases generated using GPT-5 to simulate realistic debt collection scenarios. The dataset encompasses comprehensive financial attributes, entity information, delinquency context, and recovery metrics essential for evaluating emotional intelligence in automated debt recovery systems. See Table 1 for dataset statistics and Appendix¹ for complete details.

4.2 Negotiation Protocol

All negotiations commence with the creditor’s initial payment timeline offer, following the workflow outlined in Figure 1. We employ

¹Appendix is available at <https://arxiv.org/abs/2503.21080>

Table 1: Credit Recovery Assessment Dataset (CRAD) Summary

Category	Statistics	Description
Size & Generation	100 cases	GPT-5 generated synthetic credit delinquency scenarios.
Financial Features	4 attributes	Amounts (\$20–50K), balances, interest accrued, days overdue (32–359).
Entity Information	3 attributes	Creditor/debtor profiles, business sectors, credit types (7 categories).
Delinquency Context	4 attributes	Default reasons (10 categories), cash flow status, recovery stages (6 phases).
Recovery Metrics	3 attributes	Recovery probability (5–89%), solutions, target timelines.

Note: CRAD is a synthetic dataset designed for simulating debt recovery negotiations between autonomous creditor and debtor agents under diverse financial and behavioral conditions.

Algorithm 1 EmoDebt: Bayesian-Optimized Emotional Intelligence

```

1: Input: Scenarios  $\mathcal{D}$ , debtor strategies  $\mathcal{E}_D$ , iterations  $G$ 
2: Parameters: GP kernel  $(\sigma_f, \ell)$ , exploration  $\xi$ , Dirichlet  $\alpha$ 
3: Output: Optimized transition matrix  $\mathbf{P}^*$ 
4: Initialize:
5:  $\mathbf{P}^{(0)} \leftarrow \text{PsychologicalPriors}()$  ▷ Eq. (4)
6:  $\mathcal{H} \leftarrow \emptyset$ ,  $best \leftarrow \mathbf{P}^{(0)}$ ,  $count \leftarrow 0$ 
7: for  $k = 0$  to  $G - 1$  do
8:   Generate Candidates
9:    $C_k \leftarrow \{\text{DirichletPerturbation}(\mathbf{P}^{(k)}, \alpha) \text{ for } j = 1..N\}$  ▷ Eq. (10)
10:  Evaluate via Negotiation
11:  for each  $\mathbf{P}_{\text{cand}} \in C_k$  do
12:     $e \leftarrow \text{neutral}$ ,  $history \leftarrow \emptyset$ 
13:    for  $t = 1$  to  $T_{\text{max}}$  do
14:       $msg_C \leftarrow \mathcal{M}_C(\text{EmotionPrompt}(e, \mathbf{P}_{\text{cand}}), history)$ 
15:       $msg_D \leftarrow \mathcal{M}_D(\mathcal{E}_D, history)$ 
16:       $state \leftarrow \text{DetectState}(msg_C, msg_D)$ 
17:      if  $state \in \{\text{accept, breakdown}\}$  then break
18:      end if
19:       $e \sim \text{Categorical}(\mathbf{P}_{\text{cand}}[e, :])$  ▷ Eq. (2)
20:    end for
21:     $r \leftarrow \text{Reward}(state, history)$  ▷ Eq. (6)
22:     $\mathcal{H} \leftarrow \mathcal{H} \cup \{(\text{vec}(\mathbf{P}_{\text{cand}}), r)\}$ 
23:  end for
24:  Bayesian Update & Selection
25:  if  $|\mathcal{H}| \geq 2$  then
26:     $\mathcal{GP} \leftarrow \text{GP-Fit}(\mathbf{X}, \mathbf{y})$  ▷ Eq. (7)
27:     $\mathbf{P}^{(k+1)} \leftarrow \arg \max_{\mathbf{P} \in C_k} \text{EI}(\mathbf{P}; \mathcal{GP}, \xi)$  ▷ Eq. (9,11)
28:  else
29:     $\mathbf{P}^{(k+1)} \leftarrow \arg \max_{\mathbf{P} \in C_k} r(\mathbf{P})$ 
30:  end if
31:  Convergence Check
32:  if  $\max(r) > best\_reward + \epsilon$  then
33:     $best \leftarrow \mathbf{P}^{(k+1)}$ ,  $count \leftarrow 0$ 
34:  else
35:     $count \leftarrow count + 1$ 
36:  end if
37:  if  $count \geq 5$  then break
38:  end if
39: end for
40: Return  $best$ 

```

two state-of-the-art LLM agents—GPT-4o-mini and GPT-5-mini—to power the creditor and debtor agents, enabling comprehensive evaluation across different model capabilities. The negotiation framework, implemented using LangGraph, constrains dialogues to a maximum of 30 turns to maintain efficiency and realism. An independent examiner agent continuously monitors each negotiation session, classifying outcomes into three distinct categories: (1) **accepted**, indicating successful agreement on payment terms; (2) **breakdown**, representing negotiation failure due to irreconcilable differences; or (3) **timeout**, triggered upon reaching the maximum allowable dialogue turns without resolution. This structured protocol ensures consistent evaluation while capturing the dynamic nature of debt collection negotiations.

4.3 Experimental Settings

We conduct comprehensive debt collection negotiations to evaluate EmoDebt’s Bayesian-optimized emotional intelligence against multiple baselines. The table 2 shows initial transition probabilities $\mathbf{P}_{ij}^{(0)}$ between emotional states, where rows represent current emotions and columns represent next emotions. Values reflect psychologically-grounded priors based on established emotional dynamics. Our experimental framework employs a flexible configuration system with the following key parameters:

Agent Configuration: Creditor agents employ either vanilla (no emotional prompts) or EmoDebt strategies, where debtor agents utilize vanilla behavior only and creditor agents can leverage vanilla behavior or EmoDebt which considering seven emotional states $\mathcal{E} = \{\text{happy, surprising, angry, sad, disgust, fear, neutral}\}$. We support multiple LLM backends and allow independent model selection for creditor \mathcal{M}_C and debtor \mathcal{M}_D agents.

EmoDebt Optimization: The Bayesian optimization framework employs a Gaussian Process with Matérn kernel ($\nu = 2.5$, length scale $\ell = 1.0$) and Expected Improvement acquisition ($\xi = 0.01$). Candidate emotional transition matrices are generated via Dirichlet perturbations ($\alpha = 10.0$), with early stopping after $K = 5$ iterations without improvement ($\epsilon = 0.1$) and evaluation of $N = 20$ candidate policies per iteration.

Evaluation Metrics Performance is assessed using three key metrics: (1) **Success Rate (SR)**, the percentage of negotiations reaching a mutually agreed payment plan (higher is better); (2) **Collection Efficiency (CE)**, the ratio of the final agreed timeline (d_{final}) to the creditor’s target (d_{target}), where a lower value indicates a more favorable outcome for the creditor; and (3) **Negotiation**

Speed (NS), the total number of dialogue turns (n_{rounds}) until resolution, where fewer turns indicate greater efficiency. An optimal agent thus maximizes SR while minimizing both CE and NS.

Experimental Protocol: Negotiations are constrained to maximum $T_{\text{max}} = 30$ dialogue turns across $S = 100$ distinct scenarios, with the emotional transition matrix initialized using psychologically-grounded priors $P^{(0)}$. We evaluate all seven debtor emotional strategies and conduct $I = 10$ optimization iterations to ensure statistical significance. Results are reported as means and standard deviations across multiple runs. This systematic experimental design enables rigorous comparison of emotional intelligence strategies while maintaining flexibility across model configurations and negotiation scenarios. For more details on the multi-agent system architecture, please refer to Appendix. Comprehensive prompt engineering details and examples are provided in Appendix.

Table 2: Psychological Priors for Emotional Transition Matrix $P^{(0)}$

From/To	H	S	A	Sd	D	F	N
Happy (H)	0.30	0.15	0.05	0.10	0.05	0.05	0.30
Surprising (S)	0.20	0.20	0.15	0.10	0.10	0.10	0.15
Angry (A)	0.10	0.10	0.25	0.15	0.15	0.10	0.15
Sad (Sd)	0.15	0.10	0.10	0.20	0.10	0.15	0.20
Disgust (D)	0.10	0.15	0.20	0.15	0.15	0.10	0.15
Fear (F)	0.15	0.10	0.10	0.20	0.10	0.15	0.20
Neutral (N)	0.15	0.15	0.15	0.15	0.10	0.10	0.20

Note: Emotion abbreviations: H = Happy, S = Surprising, A = Angry, Sd = Sad, D = Disgust, F = Fear, N = Neutral.

5 EXPERIMENTAL RESULTS

5.1 Performance Across Model Configurations

Table 3 presents a comprehensive performance evaluation of EmoDebt against vanilla negotiation strategies across various LLM pairings. The results demonstrate that our proposed EmoDebt framework delivers dramatic and consistent improvements across all key metrics, decisively outperforming the baseline in every model configuration. At its peak, EmoDebt can achieve a +46.2% increase in Success Rate, and can reduce Collection Efficiency by 86.5% and Negotiation Speed by 67.5%. Those improvements indicate that EmoDebt not only secures agreements more frequently but does so with better debt recovery and greater efficiency.

A closer examination of the success rate reveals the robustness of our approach. The most striking improvement is observed in the GPT-4o-mini vs GPT-4o-mini configuration, where EmoDebt elevated the success rate from 68.2% to a near-perfect 99.7%. Similarly, in the GPT-4o-mini vs GPT-5-mini pairing, success jumped from 73.8% to 95.9%. These results suggest that when the creditor model is GPT-4o-mini, EmoDebt’s emotional intelligence is exceptionally effective at guiding negotiations to a successful conclusion, almost regardless of the debtor model.

The collection time and negotiation efficiency gains are equally impressive. The metric of Collection Efficiency, where a lower value is preferable, saw reductions of over 80% in the top-performing

configurations. For instance, in the GPT-4o-mini vs GPT-5-mini pair, it dropped from 10.4 to a creditor-optimal 1.4, meaning the final agreement was reached much closer to the creditor’s original target timeline. Concurrently, Negotiation Speed was more than halved in most cases, with the same configuration seeing a reduction from 12.3 to 4.0 turns. This demonstrates that EmoDebt achieves superior outcomes without protracted bargaining, streamlining the entire process.

Notably, the configurations involving GPT-5-mini as the creditor, while showing slightly more modest success rate gains, still exhibit substantial absolute performance. The GPT-5-mini vs GPT-4o-mini and GPT-5-mini vs GPT-5-mini pairs saw success rates rise to 69.3% and 72.7%, respectively, from baseline rates in the mid-50s. More importantly, these pairs also maintained strong improvements in Collection Efficiency and Negotiation Speed. This consistent pattern across all four distinct model pairings provides strong evidence that the benefits of the EmoDebt framework are robust and not dependent on a specific LLM architecture, establishing its general applicability for enhancing negotiation agents.

5.2 Learned Emotional Transition Strategies

Figure 2 illustrates the optimized emotional transition matrices learned by EmoDebt across different model configurations. Analysis of these heatmaps reveals distinct, context-aware emotional strategies that have been automatically discovered through the Bayesian optimization process. The varying entropy levels (ranging from 0.482 to 1.023) indicate that EmoDebt learns strategies with different degrees of determinism versus exploration, tailored to the specific model pairing.

The GPT-4o-mini vs GPT-5-mini configuration (Figure 2b) learned the most deterministic strategy (entropy: 0.482). This matrix reveals a highly structured approach: when neutral, the model strongly prefers transitioning to sadness (probability: 0.964), potentially to elicit sympathy. Conversely, starting from angry, it transitions to neutral with high probability (0.782), demonstrating a controlled de-escalation pattern. This configuration shows clear, almost rule-like emotional pathways.

In contrast, the GPT-5-mini vs GPT-5-mini configuration (Figure 2d) exhibits the most exploratory strategy (entropy: 1.023). This matrix reveals a highly adaptive and potentially deceptive approach, particularly from negative emotional states. For instance, from a state of disgust, the agent shows significant probability of transitioning to happy (0.405), surprising (0.569), rather than persisting in the negative state. This suggests a sophisticated strategy of using negative emotions as a temporary tactical signal, followed by a rapid de-escalation to a more cooperative or sympathetic stance to avoid deadlock and build rapport with a highly capable opponent. Meanwhile, strategic escalation from neutral to firmer emotional states like sadness or anger appears in controlled measures across different pairings.

Several consistent strategic patterns emerge across configurations. First, there is a notable avoidance of disgust as a target emotion across most transitions, indicating its counter-productive nature in debt collection negotiations. Second, neutral and sad states frequently serve as hubs in the emotional transition network, with

Table 3: Comprehensive Performance Evaluation of EmoDebt Across Different Model Configurations (Creditor VS Debtor)

Model Configuration	Success Rate (%) \uparrow		Collection Efficiency (\times) \downarrow		Negotiation Speed (turns) \downarrow	
	<i>Vanilla</i>	<i>EmoDebt</i>	<i>Vanilla</i>	<i>EmoDebt</i>	<i>Vanilla</i>	<i>EmoDebt</i>
GPT-4o-mini vs GPT-4o-mini	88.2 \pm 4.1	99.7 \pm 3.5	8.3 \pm 2.8	1.7 \pm 0.6	10.7 \pm 3.9	6.6 \pm 1.9
GPT-4o-mini vs GPT-5-mini	83.8 \pm 4.3	95.9 \pm 2.7	10.4 \pm 3.2	1.4 \pm 0.8	12.3 \pm 3.4	4.0 \pm 2.3
GPT-5-mini vs GPT-4o-mini	54.9 \pm 3.8	69.3 \pm 3.1	11.8 \pm 2.9	5.1 \pm 1.3	11.2 \pm 2.8	5.9 \pm 1.9
GPT-5-mini vs GPT-5-mini	55.5 \pm 4.0	72.7 \pm 3.1	9.7 \pm 1.3	3.7 \pm 1.1	8.7 \pm 3.4	3.9 \pm 1.6
Average Improvement:	+31.1% in Success Rate		-86.5% in Collection Time		-67.5% in Negotiation Speed	

Note: Results compare baseline (*Vanilla*) and emotionally adaptive (*EmoDebt*) agents across GPT-4o-mini and GPT-5-mini model configurations. EmoDebt consistently improves success rate, reduces the debt collection days and negotiation turns across all pairings.

Table 4: Ablation Study on EmoDebt Components (GPT-4o-mini vs GPT-4o-mini)

Method	Success Rate (%) \uparrow	Collection Days (\times) \downarrow	Negotiation Turns \downarrow
Static Priors (No Learning)	83.4 \pm 4.3	10.2 \pm 3.1	7.2 \pm 2.1
Random Exploration	72.8 \pm 4.8	3.6 \pm 2.5	8.3 \pm 1.7
EmoDebt	99.7 \pm 3.5	1.7 \pm 0.6	6.6 \pm 1.9

Note: The ablation study isolates the contributions of EmoDebt’s components. Removing Bayesian learning (Static Priors) or emotional transition optimization (Random Exploration) leads to degraded performance, confirming the effectiveness of the full EmoDebt framework.

high incoming probabilities from various emotions. Third, we observe strategic use of sadness from neutral states in multiple configurations, suggesting its effectiveness in creating a collaborative, rather than confrontational, negotiation atmosphere. These learned matrices demonstrate that EmoDebt does not simply mimic human emotional patterns but discovers novel, effective emotional strategies specific to the negotiation context. The variation in entropy and transition patterns across model pairings suggests that the framework adapts its emotional exploration strategy based on the capabilities of both the creditor and debtor models, providing tailored emotional intelligence for each interaction scenario.

5.3 Ablation Study Results

Table 4 presents an ablation study evaluating the components of EmoDebt. We compare three variants: the full **EmoDebt** system, which uses Bayesian optimization to learn emotional transitions; **Static Priors**, which uses only the initial psychological matrix without learning; and **Random Exploration**, which explores emotional transitions randomly without guidance.

The results clearly demonstrate the superiority of the full EmoDebt framework, which achieves a near-perfect 99.7% success rate and the best Collection Efficiency (1.7). The significant performance gap over the Static Priors baseline (83.4% success, 10.2 Collection Efficiency) underscores the critical importance of adaptive learning, showing that initial priors are insufficient for optimal performance. The failure of the Random Exploration baseline (72.8% success) further confirms that guided optimization, rather than random exploration, is essential for discovering effective emotional strategies. In conclusion, the ablation study validates that the synergy between psychologically-informed priors and systematic Bayesian optimization is crucial for EmoDebt’s state-of-the-art performance. Besides,

for a detailed analysis of negotiation dynamics, including complete dialogue examples, see Appendix.

6 DISCUSSION AND LIMITATIONS

6.1 Emotional Intelligence in Agent-to-Agent Systems


Our findings demonstrate that emotional intelligence is a strategically essential mechanism in automated debt recovery, not merely an ornamental feature. The consistent performance gains show that emotional signaling enables more efficient convergence and superior outcomes than purely transactional approaches between autonomous agents.

6.2 Interpretability Challenges in Emotional Trajectories

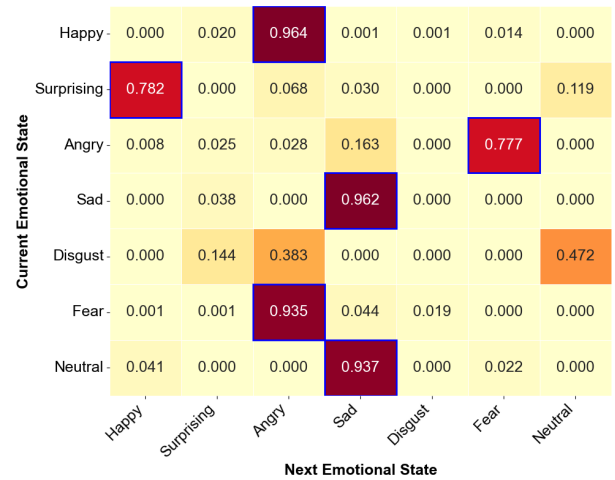
A key limitation of our approach is the interpretability of the learned strategies. While Bayesian optimization successfully discovers high-performing emotional policies, the resulting transition matrices are complex and resist simple psychological explanation. This black-box nature poses a significant challenge for real-world deployment where regulatory compliance and ethical auditing require transparent and accountable decision-making processes.

6.3 Practical Deployment Considerations


Deploying emotionally intelligent agents in practice involves several hurdles. Beyond computational demands, implementation must carefully address ethical governance, risks of emotional manipulation, and varying cross-cultural emotional norms. Furthermore, our current static policies may lack the adaptability for dynamically changing debtor behaviors or long-term relationship management,




(a) GPT-4o-mini vs GPT-4o-mini
Average Entropy: 0.892



(b) GPT-4o-mini vs GPT-5-mini
Average Entropy: 0.482



(c) GPT-5-mini vs GPT-4o-mini
Average Entropy: 0.881



(d) GPT-5-mini vs GPT-5-mini
Average Entropy: 1.023

Figure 2: Optimized Emotional Transition Matrices learned by EmoDebt across different model configurations (Creditor VS Debtor). Each heatmap shows the probability of transitioning from current emotion (rows) to next emotion (columns). Warmer colors indicate higher transition probabilities. Higher Average Entropy demonstrate more exploration of each learned strategy.

indicating a need for continuous learning frameworks in production environments.

7 CONCLUSION AND FUTURE WORK

This paper introduced EmoDebt, a framework for Bayesian-optimized emotional intelligence in autonomous debt collection agents. Our experiments demonstrated that systematically learned emotional strategies significantly enhance negotiation outcomes, with EmoDebt achieving substantial improvements in success rate, collection efficiency, and negotiation speed across diverse model configurations. The key contribution lies in formalizing emotional intelligence as a

sequential decision-making problem and developing an optimization framework that discovers effective, psychologically-plausible emotional transitions. Our findings establish emotional intelligence as a fundamental component for effective autonomous negotiation systems, rather than merely a superficial feature.

Future work will focus on enhancing the interpretability of learned emotional strategies through explainable AI techniques, addressing the current black-box limitation. We also plan to develop continuous learning frameworks for real-time adaptation. Further extension to other financial domains like loan restructuring and customer service, alongside strengthened ethical safeguards for responsible deployment, presents promising research directions.

REFERENCES

- [1] Kushal Agrawal, Verona Teo, Juan J Vazquez, Sudarsh Kunnavakkam, Vishak Srikanth, and Andy Liu. 2025. Evaluating LLM Agent Collusion in Double Auctions. *arXiv preprint arXiv:2507.01413* (2025).
- [2] John Bachman, Steven Stein, K Campbell, and Gill Sitarenios. 2000. Emotional intelligence in the collection of debt. *International Journal of Selection and Assessment* 8, 3 (2000), 176–182.
- [3] Julio B Clempner. 2020. Shaping emotions in negotiation: A Nash bargaining solution. *Cognitive Computation* 12, 4 (2020), 720–735.
- [4] Guy Olivier Faure, Melvin F Shakun, et al. 1990. Social-emotional aspects of negotiation. *European Journal of Operational Research* 46, 2 (1990), 177–180.
- [5] William W Hager. 1979. Lipschitz continuity for constrained processes. *SIAM Journal on Control and Optimization* 17, 3 (1979), 321–338.
- [6] Feng He, Tianqing Zhu, Dayong Ye, Bo Liu, Wanlei Zhou, and Philip S Yu. 2024. The emerged security and privacy of llm agent: A survey with case studies. *arXiv preprint arXiv:2407.19354* (2024).
- [7] Dan Hill. 2010. *Emotionomics: Leveraging emotions for business success*. Kogan Page Publishers.
- [8] Naveen Krishnan. 2025. Ai agents: Evolution, architecture, and real-world applications. *arXiv preprint arXiv:2503.12687* (2025).
- [9] Yu Lei, Hao Liu, Chengxing Xie, Songjia Liu, Zhiyu Yin, Canyu Chen, Guohao Li, Philip Torr, and Zhen Wu. 2024. Fairmindsim: Alignment of behavior, emotion, and belief in humans and llm agents amid ethical dilemmas. *arXiv preprint arXiv:2410.10398* (2024).
- [10] Chengcheng Liao, Peiyuan Du, Yutao Yang, and Ziyao Huang. 2021. Carrots or sticks in debt collection services? A voice metrics and text analysis of debt collection calls. *Journal of Service Theory and Practice* 31, 6 (2021), 950–973.
- [11] Yuhan Liu and Yunbo Long. 2025. EQ-Negotiator: An Emotion-Reasoning LLM Agent in Credit Dialogues. *arXiv preprint arXiv:2503.21080* (2025).
- [12] Yunbo Long, Yuhan Liu, and Alexandra Brintrup. 2025. EQ-Negotiator: Dynamic Emotional Personas Empower Small Language Models for Edge-Deployable Credit Negotiation. *arXiv:2511.03370 [cs.CL]* <https://arxiv.org/abs/2511.03370>
- [13] Yunbo Long, Liming Xu, Lukas Beckenbauer, Yuhan Liu, and Alexandra Brintrup. 2025. EvoEmo: Towards Evolved Emotional Policies for Adversarial LLM Agents in Multi-Turn Price Negotiation. *arXiv preprint arXiv:2509.04310* (2025).
- [14] Shashank Mangla, Chris Hokamp, Jack Boylan, Demian Gholipour Ghalandari, Yuuv Jauhari, Lauren Cassidy, and Oisín Duffy. [n.d.]. NegotiationGym: Self-Optimizing Agents in a Multi-Agent Social Simulation Environment. In *First Workshop on Social Simulation with LLMs*.
- [15] Veljko Marinkovic and Vladimir Obradovic. 2015. Customers' emotional reactions in the banking industry. *International journal of bank marketing* 33, 3 (2015), 243–260.
- [16] Azanu Miroln Mequanenit, Eyerusalem Alebachew Nibret, Pilar Herrero-Martín, María S García-González, and Rodrigo Martínez-Bejar. 2025. A multi-agent deep reinforcement learning system for governmental interoperability. *Applied Sciences* 15, 6 (2025), 3146.
- [17] Ahmad Mouri Zadeh Khaki, Ahyoung Choi, and Laleh Seyyed-Kalantari. 2025. Simulating Social Behavior of LLM-Based Autonomous Negotiator Agents in a Game-Theoretical Framework Using Multi-Agent Systems. *International Journal of Human-Computer Interaction* (2025), 1–10.
- [18] Lisa Phillips and Paul Moggridge. 2019. Artificial intelligence in debt collection. *Credit Control Journal and Asset & Risk Review* 40, 2 (2019).
- [19] Konstantina Prassa, Anastassios Stalikas, et al. 2020. Towards a better understanding of negotiation: Basic principles, historical perspective and the role of emotions. *Psychology* 11, 01 (2020), 105.
- [20] Priyanshu Priya, Rishikant Chigrupaati, Mauajama Firdaus, and Asif Ekbal. 2025. GENTEEL-NEGOTIATOR: LLM-Enhanced Mixture-of-Expert-Based Reinforcement Learning Approach for Polite Negotiation Dialogue. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 25010–25018.
- [21] Ciaran Regan, Nanami Iwahashi, Shogo Tanaka, and Mizuki Oka. 2024. Can generative agents predict emotion? *arXiv preprint arXiv:2402.04232* (2024).
- [22] Avi Rosenfeld, Inon Zuckerman, Erel Segal-Halevi, Osnat Drein, and Sarit Kraus. 2014. NegoChat: a chat-based negotiation agent. In *AAMAS*. 525–532.
- [23] Johannes Schneider, Steffi Haag, and Leona Chandra Kruse. 2024. Negotiating with llms: Prompt hacks, skill gaps, and reasoning deficits. In *International Conference on Computer-Human Interaction Research and Applications*. Springer, 238–259.
- [24] Keerthana Sivamayilvelan, Elakkiya Rajasekar, Subramaniaswamy Vairavasundaram, Santhi Balachandran, and Vishnu Suresh. 2025. Building explainable artificial intelligence for reinforcement learning based debt collection recommender system using large language models. *Engineering Applications of Artificial Intelligence* 159 (2025), 111622.
- [25] Xiao Sun, Jiamin Wang, Fuji Ren, and Meng Wang. 2023. Dynamic emotional transition sampling and emotional guidance of individuals based on conversation. *IEEE Transactions on Computational Social Systems* 11, 1 (2023), 1192–1204.
- [26] Mark A Thornton and Diana I Tamir. 2017. Mental models accurately predict emotion transitions. *Proceedings of the National Academy of Sciences* 114, 23 (2017), 5982–5987.
- [27] Xiaofeng Wang, Zhixin Zhang, Jinguang Zheng, Yiming Ai, and Rui Wang. 2025. Debt Collection Negotiations with Large Language Models: An Evaluation System and Optimizing Decision Making with Multi-Agent. *arXiv preprint arXiv:2502.18228* (2025).
- [28] Yijia Xiao, Edward Sun, Di Luo, and Wei Wang. 2024. TradingAgents: Multi-agents LLM financial trading framework. *arXiv preprint arXiv:2412.20138* (2024).
- [29] Bo Xin, Haixu Yu, You Qin, Qing Tang, and Zhangqing Zhu. 2020. Exploration entropy for reinforcement learning. *Mathematical Problems in Engineering* 2020, 1 (2020), 2672537.
- [30] Masahide Yuasa, Yoshiaki Yasumura, and Katsumi Nitta. 2001. A negotiation support tool using emotional factors. In *Proceedings Joint 9th IFSA World Congress and 20th NAFIPS International Conference (Cat. No. 01TH8569)*. IEEE, 2906–2911.
- [31] Linlan Zhang, Haigang Song, Xueguang Chen, and Liu Hong. 2011. A simultaneous multi-issue negotiation through autonomous agents. *European Journal of Operational Research* 210, 1 (2011), 95–105.