

What Do You Want From Me?

Adapting Systems to the Uncertainty of Human Preferences

Carlos Gavidia-Calderon
Amel Bennaceur
carlos.gavidia-calderon@open.ac.uk
amel.bennaceur@open.ac.uk
The Open University
United Kingdom

Anastasia Kordoni
Mark Levine
a.kordoni@lancaster.ac.uk
mark.levine@lancaster.ac.uk
Lancaster University
United Kingdom

Bashar Nuseibeh
bashar.nuseibeh@open.ac.uk
The Open University
United Kingdom
Lero - University of Limerick
Ireland

ABSTRACT

Autonomous systems, like drones and self-driving cars, are becoming part of our daily lives. Multiple people interact with them, each with their own expectations regarding system behaviour. To adapt system behaviour to human preferences, we propose and explore a game-theoretic approach. In our architecture, autonomous systems use sensor data to build game-theoretic models of their interaction with humans. In these models, we represent human preferences with *types* and a probability distribution over them. Game-theoretic analysis then outputs a strategy, that determines how the system should act to maximise utility, given its beliefs over human types. We showcase our approach in a search-and-rescue (SAR) scenario, with a robot in charge of locating victims. According to social psychology, depending on their identity some people are keen to help others, while some prioritise their personal safety. These social identities define what a person favours, so we can map them directly to game-theoretic types. We show that our approach enables a SAR robot to take advantage of human collaboration, outperforming non-adaptive configurations in average number of successful evacuations.

CCS CONCEPTS

• **Computer systems organization** → **Robotics**; • **Human-centered computing** → **Collaborative interaction**.

KEYWORDS

autonomous systems, social identity, game theory

ACM Reference Format:

Carlos Gavidia-Calderon, Amel Bennaceur, Anastasia Kordoni, Mark Levine, and Bashar Nuseibeh. 2022. What Do You Want From Me? Adapting Systems to the Uncertainty of Human Preferences. In *New Ideas and Emerging Results (ICSE-NIER'22)*, May 21–29, 2022, Pittsburgh, PA, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3510455.3512791>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ICSE-NIER'22, May 21–29, 2022, Pittsburgh, PA, USA

© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-9224-2/22/05...\$15.00
<https://doi.org/10.1145/3510455.3512791>

1 INTRODUCTION

Autonomous systems are increasingly operating among us. We share roads with self-driving cars [4], and drones deliver parcels at our doorsteps [15]. These systems need to adapt to the complexity of the real world. For instance, changes in weather conditions require adaptation from delivery drones, since harsh weather can distort sensor data needed for object detection [3]. These systems should also adapt to the needs, expectations, and preferences of the human beings they interact with [9]. For example, self-driving car users differ on preferred speed and route to destination. This depends on their sense of urgency, how safe they feel, and even the trust they have in the system. Recognising and adapting autonomous systems to the differences among humans is key for enabling robot-human cooperation. More importantly, if a system does not align with human expectations, it is less likely to be trusted and used [14].

Dealing with uncertain human preferences is not only a problem for autonomous systems. Game theory, the study of mathematical models of conflict and cooperation [12], is also interested in scenarios where interacting agents are uncertain about each other's preferences. Game theory studies scenarios — called games — where rational and self-interested agents interact, and these interactions affect the utility they perceive. This definition suits board games, card games, markets, and even software development teams [6]. In games like chess, agent preferences are unambiguous: they both have full visibility of the board and want to win the game. In contrast, in games like poker agent preferences are uncertain, since they have no access to the deck or their opponent's hand. *Bayesian games* are game-theory's approach to model such scenarios.

One of the most representative examples of Bayesian games are sealed-bid auctions [13]. Here agents are the bidders and, although they know their own valuation of the auctioned item, they ignore the valuation of competing bidders. An agent's private information, like the item valuation, is called its *type*. In a Bayesian game, the probability distribution over types is common knowledge. An agent's strategy in a Bayesian game defines their behaviour according to their type. At equilibrium, the strategy of every agent is the best response to the strategies of the other agents. An equilibrium is stable — deviation from equilibrium would result in utility loss — so we expect agents to adopt these strategies. In fact, there is empirical evidence that experienced bidders perform equilibrium strategies when engaging in internet auctions [19].

In this paper, we propose autonomous systems equipped with Bayesian game models of human-system interactions (section 2). Using sensors, autonomous systems detect when they engage in an interaction that requires adaptation. When this happens, the

system uses sensor data to complete a game-theoretic model. Finally, a game solver analyses the model to obtain the strategies at equilibrium. The autonomous system then translates this strategy to actuator commands.

The game-theoretic models behind system adaptation are, in essence, models of social behaviour. Thus, designing these models can benefit from psychological theories of group behaviour. In section 3, we use the social identity approach to design a search-and-rescue (SAR) robot. The robot’s mission is to traverse a disaster zone and search for victims. Once the robot locates a victim, it can either guide the victim to safety, or ask for support from first-responders in the area. Social psychology suggests that, in an emergency, most affected people develop a group identity favouring cooperation [10]. To reflect this, our SAR robot builds a game model of human interaction with two types. Survivors with the *group-identity type* prefer to help victims. If the robot offers to guide them to safety, it can count on surrounding victims being assisted. In contrast, if the robot faces a survivor with the *personal-identity type*, it would be better to request first-responder support, to avoid victims being abandoned.

We compare our adaptive robot with a proself-oriented and a prosocial-oriented one (section 4). The proself-oriented robot always requests first-responder support, assuming other survivors will not assist victims. The prosocial-oriented robot assumes cooperation will always happen, so it always offers guidance. The adaptive robot performs the best regarding average successful evacuations, suggesting the viability of our proposal.

In this work, we therefore show how game-theory and social psychology can enable autonomous systems’ adaption to uncertain human preferences. Like the SAR scenario, we believe there are other situations that can also benefit from our approach.

2 ADAPTING TO HUMAN PREFERENCES

Our goal is to build systems that, during operation, autonomously optimise their response to the preferences of the humans they engage. To accomplish this, we equip systems with *autonomic managers* that control their behaviour via sensors and actuators. Following the MAPE-K model [8], the autonomic manager has five responsibilities:

- (1) The manager *monitors* sensor data to detect interactions with humans that require adaptation.
- (2) The *knowledge* is represented by a game-theoretic model of the interaction. The manager updates this model frequently, according to the person the system is facing.
- (3) Each time the model is updated, the manager *analyses* the model by calculating an equilibrium strategy for the managed system.
- (4) During *planning* the manager translates this strategy into actuator commands.
- (5) Finally, the system’s actuators *execute* these commands, completing the interaction with the human.

As shown in Figure 1, we distribute these responsibilities among four components. At runtime, the Game Selector notifies the Game Builder that an interaction of type G' is taking place, given sensor readings X . The Game Builder then produces a game-theoretic model G of such interaction. G takes the form of a game tree, like the

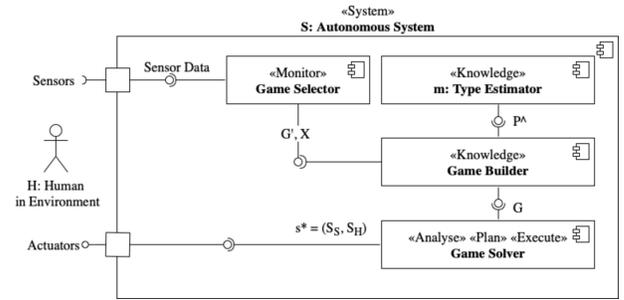


Figure 1: Component diagram of our adaptive architecture.

one in Figure 2. While some elements of G must be obtained at runtime —like the probabilities \hat{P} from the Type Estimator— engineers can incorporate other elements at design time. For instance, in our SAR scenario (section 3) the agents, actions, and utility functions do not depend on runtime information. Finally, the Game Solver solves G obtaining the equilibrium strategy S_S for the system, that can be translated to actuator commands. In the rest of this section, we describe each component in detail.

Game Selector. During operation, an autonomous system can interact with multiple people. Not every interaction needs an adaptation to the person’s preferences. For example, a self-driving car does not need to adapt to every person it encounters while waiting at a zebra crossing. In our architecture, the system supports a set of interactions that it can model for adaptation purposes. The role of the Game Selector component is to detect, based on sensor data, if the autonomous system is engaging in an interaction that requires adaptation. When this happens, it forwards the necessary information to the Game Builder component so it can model such interaction using game theory.

Game Builder. As discussed in section 1, Bayesian games are an adequate representation for handling uncertain human preferences. We represent Bayesian games using game trees. Per interaction, the Game Builder produces a game model G with the following elements:

- The agents S (autonomous system) and H (human) engaging in the interaction.
- A_S are the actions the system can perform, and A_H the actions the human can perform during the interaction.
- In the game tree, the system performs actions at the nodes belonging to N_S . The human acts at the nodes belonging to N_H .
- T is the type space for the human agent. Each type $t \in T$ is associated to a node $n \in N_S$. These nodes have the same available actions.
- $P : T \mapsto [0, 1]$ is a probability distribution over the human’s type space.
- The interaction ends when reaching terminal nodes in N_E .
- The function $u_S : N_E \mapsto \mathbb{R}$ produces the system’s utility when reaching a node in N_E . The function $u_H : N_E \times T \mapsto \mathbb{R}$ does the same for the human agent, given their utility also depend on their type.

- The function $\chi_S : N_S \mapsto 2^{A_S}$ selects which actions from A_S are available at node $h \in N_S$. The corresponding function for the human agent is $\chi_H : N_H \times T \mapsto 2^{A_H}$.
- The function $\sigma_S : N_S \times A_S \mapsto N_H \cup N_E$ selects a human's node $h \in N_H$ following a system's action. The function $\sigma_H : N_H \times A_H \times T \mapsto N_S \cup N_E$ does node selection after a human's action.

In the context of the general Bayesian game framework, the game G has additional restrictions: G is limited to two agents, with multiple types only for the person agent. Given that the number of players and type spaces have a large impact on model size, we adopt these model constraints to keep the model tractable. Also, there is the fact that a system can require multiple interaction models. For example, a self-driving car interacts with the driver, the passengers, and police officers, at different times. Each one of these interactions, depending on context, might need their own game-theoretic model.

Type Estimator. The probability distribution over types P is a key element of our models. Without it, the game G becomes a *game of incomplete information* with very limited applications besides worst-case scenario analysis [7]. Also, P must be obtained at runtime — each person encountered during operation has different probability values — so its calculation must be efficient. Given its importance, we have a dedicated component for obtaining P in our architecture.

Based on a relevant sensor data, the Type Estimator component computes \hat{P} , an estimate of the type probabilities P for the person interacting with the autonomous system, given a game G . The component uses the model $m_\theta : T \times \mathbb{R}^n \mapsto [0, 1]$ to compute probability values for each type $t \in T$ and a vector X of n sensor readings.

Game Solver. The Game Builder component produces a game-theoretic model G for a given human interaction with the system. The Game Solver component obtains the *equilibrium profiles* of that model. For a game G , a profile contains two strategies: a strategy s_S for the autonomous system and a strategy s_H for the human agent. As mentioned in section 1, at equilibrium s_S is the system's best response to the person's strategy s_H , and vice versa.

For an autonomous system, the strategy s_S assigns a probability to each action in $\chi_S(n)$ for each node $n \in N_S$. The human's strategy s_H depends on their type: it assigns a probability to each action in $\chi_H(n, t)$ for every node in $n \in N_H$ and type $t \in T$.

There are multiple algorithms for obtaining equilibria [16], and models can have more than one equilibrium profile. It is the Game Solver's responsibility to transform the equilibrium profiles into actionable commands for the system's actuators. For example, in a self-driving car the system's strategy in a profile can take the form of a target driving speed, or a neighbourhood to avoid. The Game Solver needs to translate these generic statements into specific commands for the car's navigation system.

3 A SEARCH-AND-RESCUE ROBOT

Emergency services actively use robots for rescue operations [2]. Among their many advantages, they can traverse the disaster area without exposing first-responders to danger. In our scenario, we deploy an autonomous search-and-rescue robot at a disaster zone. Its job is to search for victims and secure their evacuation. At the

disaster zone, the robot can encounter affected individuals, victims with injuries, and first-responders coordinating the SAR operations. The robot has two ways to facilitate victim safe evacuation: 1) after locating a victim, it can guide them to safety using its navigation capabilities; or 2) it can request assistance from a human first-responder via radio, informing their location. However, in most emergency situations, first-responders are scarce, busy, and sometimes unable to access the affected area.

The SAR robot supports adapting to a single interaction: finding two people in close proximity, a *victim* that is unable to move and a *survivor* that is healthy. The rescue team wants to maximise the number of successful evacuations, so in this scenario they prefer that the robot guides both people to a secure location. By not requesting first-responder support, they can invest their efforts in evacuating additional victims. But, for both people to evacuate under robot guidance, we need help from the survivor. Otherwise, the victim would be left behind.

For our SAR robot to adapt using Bayesian game models, we need types that reflect the likelihood for a person in the area to help. We rely on social psychology and the social identity framework [18] to inform type selection. Social psychology research establishes that identity determines people behaviour. People may act upon a *personal identity*, defined by their personality; or a social identity, defined by their group membership. This *group identity* can arise from being a woman, a student, or even an engineer. Each identity comes with beliefs, norms, and values that affect human behaviour [17]. In an emergency context, people acting upon their group identity are more likely to exhibit pro-social behaviours, like helping others [10]. In contrast, people acting upon their personal identities are more likely to show self-interest [5].

When the robot's Game Selector component detects this interaction, it forwards sensor data to the Game Builder component to generate a game tree, like the one in Figure 2. This interaction model has the following elements:

- The agents are the SAR robot (S) and the survivor (H).
- $A_S = \{N, C\}$ where N corresponds to navigate and guide them to safety and C to calling for first-responder support.
- $A_H = \{N_+, N_-, C_+, C_-\}$. When the robot offers guidance (N), the survivor can help the victim (N_+) or follow the robot by themselves (N_-). In case the robot calls a first-responder (C), the survivor can wait with the victim (C_+) or evacuate by themselves without support (C_-).
- $T = \{G, P\}$, where G corresponds to the group-identity type and P to the personal-identity type.
- u_S is the expected number of successful evacuations.
- u_H is an identity metric.

The values for u_S and u_H in the tree at Figure 2 are not backed by empirical evidence, and their sole purpose is to demonstrate our approach. For example, in the bottom-left node, the robot offers guidance (N) to a survivor with group-identity type G , and they respond by accepting the offer taking the victim with them (N_+). In this case, the robot's utility is $u_S = 3$, reflecting the evacuation of 3 people: the survivor, the victim, and another survivor rescued by the first-responder. For the survivor, their utility is $u_H = 1.3$. This value is the sum of 1 for a successful evacuation plus 0.3 for following identity expectations [1].

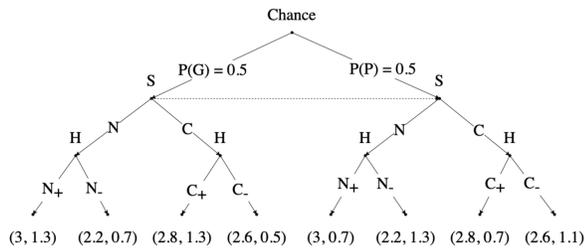


Figure 2: Game tree for the SAR scenario.

Table 1: Preliminary evaluation results.

Robot	$r_G = 50\%$	$r_G = 60\%$	$r_G = 70\%$	$r_G = 80\%$
Adaptive	91.6	92.8	93.9	95.3
Proself	89.2	89.9	90.5	91.2
Prosocial	86.1	89.0	91.6	94.0

To complete the model, we need the probability estimates \hat{P} for the group-identity G and personal-identity P types. People express their identity externally via *social identity markers*, like clothing, language, and even behaviours. We equip the Type Estimator component with a model m that translates social identity markers – inferred from sensors at the Game Selector component – into a probability distribution over the social identities in T . Given that $|T| = 2$, we use a multi-layer perceptron with a single output neuron using a sigmoid activation function. Training m should be done offline, before robot deployment.

Once the model is complete, we use Gambit [11] to obtain the strategies at equilibrium. We only consider equilibria with *subgame perfect profiles with pure strategies*, to reflect rational players with deterministic strategies. A pure strategy for the robot is to either offer guidance (N) or ask for first-responder support (C). In case Gambit finds more than one equilibrium, the robot acts conservatively by selecting C . The Game Solver component then translates the resulting strategy to commands to the robot’s navigation system for strategy N , or to GPS and radio for strategy C .

4 PRELIMINARY EVALUATION

We compare our adaptive SAR robot (section 3) against two non-adaptive ones: a *proself-oriented robot* that always requests first-responder support (C), assuming the survivor would abandon the victim; and a *prosocial-oriented robot* that always offers guidance (N) counting on the help from other survivors.

We expose each robot to 30 scenarios, having each scenario composed by 33 interactions. An interaction is the situation described in section 3 – and represented in Figure 2 – where a SAR robot encounters a survivor and a victim. During an interaction, the robot under evaluation adopts a strategy (N or C), and the survivor responds maximizing their utility u_H according to their type (G or P). The game outcome corresponds to a node $n \in N_E$, assigning a utility to the robot u_S and the survivor u_H . For example, let us consider a scenario interaction between a proself-oriented robot and a survivor with group-identity. According to its strategy C , the robot asks for first-responder support. As shown in Figure 2,

to maximise their utility the group-identity survivor should wait with the victim (C_+), resulting in an outcome with $u_S = 2.8$ and $u_H = 1.3$.

Unlike the prosocial-oriented and proself-oriented robots, the adaptive SAR robot needs sensor data to select a strategy. Using the *make_classification* method from the *Scikit-learn* library, we generate an artificial dataset with 10,000 samples representing sensor readings. Each sample has 100 features, but only 3 of them are informative. To the best of our knowledge, there is no behavioural evidence on the magnitude of group identity propagation. Hence, in our evaluation we use 50%, 60%, 70%, and 80% as plausible values for r_G , the proportion of generated samples associated to the group-identity type.

We split the dataset in two parts: 33% was used for evaluation and 67% for training the neural network m on the Type Estimator. We undersample the majority type during training to handle class imbalance. To avoid overfitting, our training process uses early stopping, with patience of 20 epochs without improvement.

Table 1 shows the evaluation results¹. For every value of r_G , the adaptive SAR robot has a higher value for \bar{u}_S , the average robot utility per scenario. We also observed that with higher values of r_G , the accuracy of m decreases due to less samples for the minority class. At $r_G = 50\%$, m has an accuracy of 89% on the validation dataset. At $r_G = 80\%$ this value is 67%. Even with this modest accuracy, the adaptive robot is able to capitalise on the emergent prosocial behaviour, outperforming the proself-oriented robot by 4.5% and the prosocial-oriented one by 1.38% when $r_G = 80\%$.

5 DISCUSSION

Game theoretic models of human-robot interaction are at the core of our approach. Model elements related to the robot, like its actions A_S and utility function u_S can be considered from an engineering perspective. However, the model elements related to human behaviour are more suited to a social psychology perspective. There is a natural correspondence between game-theoretic types and social identities. In our SAR scenario, the social identity framework provided useful abstractions for grouping people according to outcome preferences. People’s reactions to robot actions A_H and the utility they perceive u_H should be based on social psychology research.

The Type Estimator m plays a crucial role in our architecture. Inferring a person’s type T from sensor data can be complex. A person’s type can depend on features like their clothing, spoken language, or proximity to other people. Hence, designing m requires expertise in multiple computing disciplines.

We have thus proposed a novel approach for designing autonomous systems that adapt to human preferences. We suggest that it can be applied to other scenarios, like human-robot trust. This would require mapping trust levels to types, and adapting robot actions to the perceived trust level. We believe our approach is a step towards more resilient autonomous systems.

ACKNOWLEDGMENTS

This work was supported by the Engineering and Physical Sciences Research Council [grant numbers EP/V026747/1, EP/R013144/1]; and Science Foundation Ireland [grant number 13/RC/2094_P2].

¹Code available at <https://github.com/cptanalatriste/wdywfm-adaptive-robot>

REFERENCES

- [1] George A. Akerlof and Rachel E. Kranton. 2000. Economics and Identity*. *Quarterly Journal of Economics* 115, 3 (aug 2000), 715–753. <https://doi.org/10.1162/003355300554881>
- [2] H. Levent Akin, Nobuhiro Ito, Adam Jacoff, Alexander Kleiner, Johannes Pellenz, and Arnoud Visser. 2012. RoboCup Rescue Robot and Simulation Leagues. *AI Magazine* 34, 1 (dec 2012), 78. <https://doi.org/10.1609/aimag.v34i1.2458>
- [3] Mario Bijelic, Tobias Gruber, Fahim Mannan, Florian Kraus, Werner Ritter, Klaus Dietmayer, and Felix Heide. 2020. Seeing Through Fog Without Seeing Fog: Deep Multimodal Sensor Fusion in Unseen Adverse Weather. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*. Computer Vision Foundation / IEEE, 11679–11689. <https://doi.org/10.1109/CVPR42600.2020.01170>
- [4] Barry Brown. 2017. The Social Life of Autonomous Cars. *Computer* 50, 2 (feb 2017), 92–96. <https://doi.org/10.1109/MC.2017.59>
- [5] Mark Conner and Christopher J. Armitage. 1998. Extending the Theory of Planned Behavior: A Review and Avenues for Further Research. *Journal of Applied Social Psychology* 28, 15 (1998), 1429–1464. <https://doi.org/10.1111/j.1559-1816.1998.tb01685.x>
- [6] Carlos Gavidia-Calderon, Federica Sarro, Mark Harman, and Earl T. Barr. 2019. The Assessor's Dilemma: Improving Bug Repair via Empirical Game Theory. *IEEE Transactions on Software Engineering* (2019). <https://doi.org/10.1109/tse.2019.2944608>
- [7] A.R. Karlin and Y. Peres. 2017. *Game Theory, Alive*. American Mathematical Society.
- [8] Jeffrey O. Kephart and David M. Chess. 2003. The Vision of Autonomic Computing. *Computer* 36, 1 (2003), 41–50. <https://doi.org/10.1109/MC.2003.1160055>
- [9] Hadas Kress-Gazit, Kerstin Eder, Guy Hoffman, Henny Admoni, Brenna Argall, Rüdiger Ehlers, Christoffer Heckman, Nils Jansen, Ross A. Knepper, Jan Kretínský, Shelly Levy-Tzedek, Jamy Li, Todd D. Murphey, Laurel D. Riek, and Dorsa Sadigh. 2021. Formalizing and guaranteeing human-robot interaction. *Commun. ACM* 64, 9 (2021), 78–84. <https://doi.org/10.1145/3433637>
- [10] Mark Levine and Rachel Manning. 2013. Social identity, group processes, and helping in emergencies. *European Review of Social Psychology* 24, 1 (dec 2013), 225–251. <https://doi.org/10.1080/10463283.2014.892318>
- [11] Richard D. McKelvey, Andrew M. McLennan, and Theodore L. Turocy. 2021. Gambit: Software Tools for Game Theory, Version 15. <http://www.gambit-project.org>
- [12] Roger B. Myerson. 1997. *Game theory - Analysis of Conflict*. Harvard University Press. <http://www.hup.harvard.edu/catalog/MYEGAM.html>
- [13] Itzhak Rasooly and Carlos Gavidia-Calderon. 2020. The importance of being discrete: on the inaccuracy of continuous approximations in auction theory. *arXiv preprint arXiv:2006.03016* (2020).
- [14] Paul Robinette, Ayanna M. Howard, and Alan R. Wagner. 2017. Effect of Robot Performance on Human-Robot Trust in Time-Critical Situations. *IEEE Trans. Hum. Mach. Syst.* 47, 4 (2017), 425–436. <https://doi.org/10.1109/THMS.2017.2648849>
- [15] David Schneider. 2020. The delivery drones are coming. *IEEE Spectrum* 57, 1 (jan 2020), 28–29. <https://doi.org/10.1109/MSPEC.2020.8946304>
- [16] Yoav Shoham and Kevin Leyton-Brown. 2009. *Multiagent Systems - Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press.
- [17] Russell Spears. 2021. Social Influence and Group Identity. *Annual Review of Psychology* 72, 1 (2021), 367–390. <https://doi.org/10.1146/annurev-psych-070620-111818>
- [18] Henri Tajfel, John C Turner, William G Austin, and Stephen Worchel. 1979. An integrative theory of intergroup conflict. *Organizational identity: A reader* 56, 65 (1979), 9780203505984–16.
- [19] Ronald T Wilcox. 2000. Experts and amateurs: The role of experience in Internet auctions. *Marketing Letters* 11, 4 (2000), 363–374.