# Pedestrian Interaction Modelling: Leveraging Trajectory Prediction for Belief Representation

MSc Thesis

T.H.Weinans

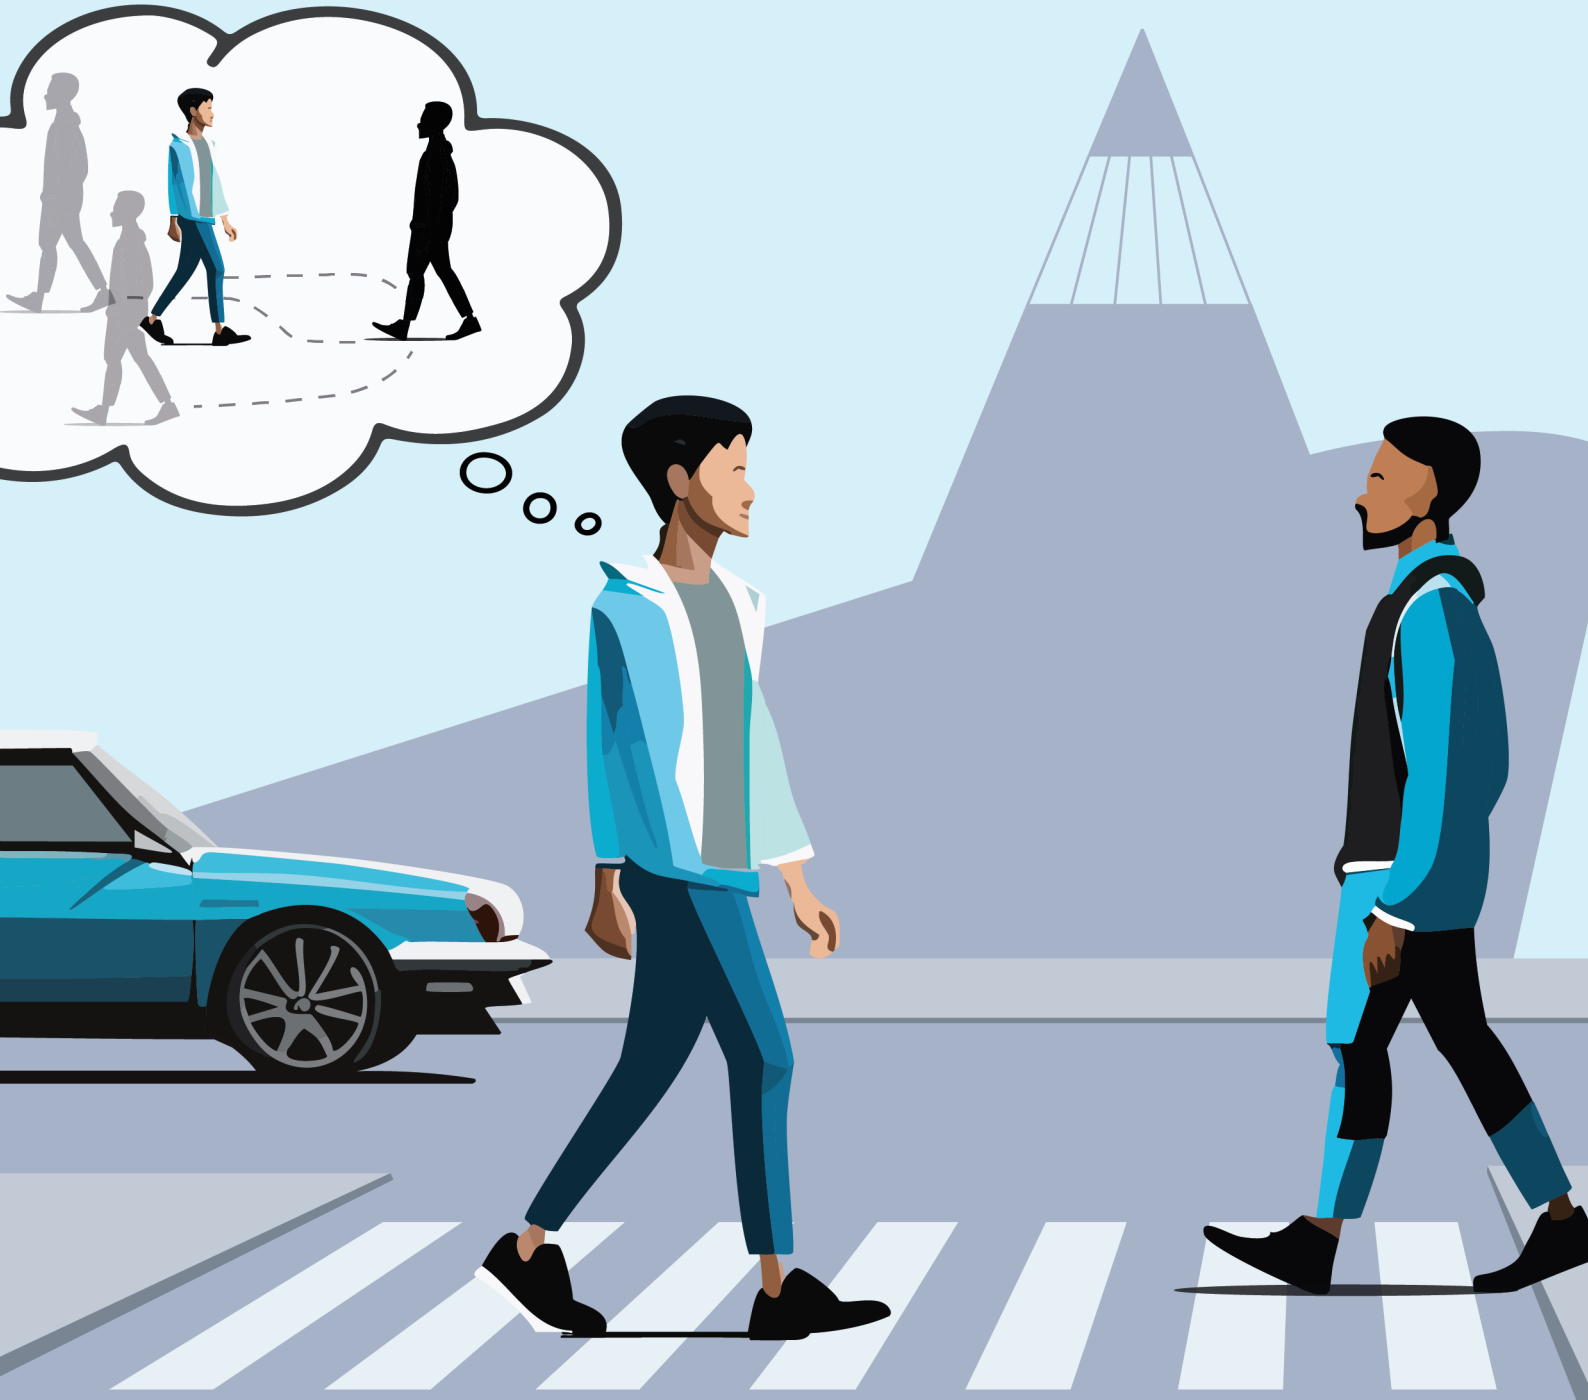

TU Delft

# Pedestrian Interaction Modelling: Leveraging Trajectory Prediction for Belief Representation

## MSc Thesis

by

T.H. Weinans

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Friday, February 2, 2024, at 09:30 AM.

Student number: 4445449
Project duration: May 3, 2023 – February 2, 2024
Thesis committee:

| | |
|---|---|
| Dr. A. Zgonnikov , | TU Delft, supervisor |
| Ir. O. Siebinga, | TU Delft, supervisor |
| Ir. A. Mészáros, | TU Delft, supervisor |
| Dr. J. Kober, | TU Delft, external member |



TU Delft

# Preface

Dear reader, in front of you lies my master thesis: “Pedestrian Interaction Modelling: Leveraging Trajectory Prediction for Belief Representation”. I have dedicated the majority of the past year to this work, which will conclude my graduation project as well as my master in Robotics at the Delft University of Technology. During this time I have immersed myself in the world of interaction modelling and trajectory prediction with the goal of making a meaningful contribution to the field of pedestrian interaction modelling. With this goal in mind, I have researched the effectiveness of integrating a probabilistic trajectory forecasting model, trained on a belief-based dataset, into an interaction modelling framework driven by communication and beliefs. To test its abilities I have created a two-dimensional simulation setup capable of simulating several interactive pedestrian traffic scenarios, of which all my findings are reported in this thesis document.

This has been the biggest project I have ever worked on and I have enjoyed it a lot. I could not have done it alone, which is why I want to thank my daily supervisors Anna and Olger, for all the feedback, brainstorming and other support you have provided throughout the entire process. Your guidance has really helped me stay motivated and focused on the goal. I also want to thank my supervisor Arkady for your expertise, encouragement, and valuable insights. I would also like to thank my family, friends, and anyone close to me who has helped me to get to this point, and finally I would like to thank you, the reader for taking the time to read my work. I hope you will enjoy it.

*T.H. Weinans*
*Delft, January 2024*

# Contents

# I

# Scientific Paper

# Pedestrian Interaction Modelling: Leveraging Trajectory Prediction for Belief Representation

T.H. Weinans
*Delft University of Technology*
Delft, The Netherlands
T.H.Weinans@student.tudelft.nl

***Abstract*—Autonomy in traffic (e.g., autonomous vehicles) could potentially benefit mobility, safety, accessibility and sustainability. However, the realisation of these advancements is highly dependent on how effective these autonomous vehicles interact with vulnerable road users such as pedestrians. Before we can understand how pedestrians will interact with autonomous vehicles, it is essential to understand how pedestrians interact among themselves in interactive traffic scenarios. Previous studies have focused on describing these scenarios with probabilistic trajectory prediction methods such as TrajFlow. However, these approaches often fall short in capturing the nuances of mutual interactions. Simple interaction models have been proposed that can describe these interactions, but neglect the influence of another person's intentions. To address this issue, in existing work the Communication-Enabled-Interaction (CEI) framework was proposed that describes interactions by modelling communication and a belief of another person's intentions. The idea of using beliefs in interaction modelling is based on the concept that people have a general but uncertain idea about the plans of other people. These beliefs are one of the fundamental aspects of the CEI framework and must therefore contain valuable information about possible decisions. That is why this study investigates the use of the probabilistic trajectory prediction method TrajFlow for the belief construction of the CEI framework. TrajFlow is trained on the belief-based Forking Paths dataset, integrated into the CEI framework, and tested in four simulated pedestrian interaction scenarios. The analysis shows that the framework is able to simulate plausible interaction behaviour, dealing with conflicting goals and trajectories in multiple simulations. By doing so, this study takes a positive step towards modelling pedestrian interactions and contributes to the broader goal of realising the benefits linked to autonomy in traffic.**

***Index Terms*—Pedestrian behaviour, interaction modelling, trajectory prediction**

## I. Introduction

In this day and age, there is a great deal of research being done towards autonomous vehicles (AVs) in traffic situations, such as self-driving cars and delivery robots. There are many promised benefits related to mobility, safety, accessibility and sustainability [1]. Despite all these anticipated benefits, the realisation of these advancements is highly dependent on how effective their interaction is with human road users, particularly vulnerable road users like pedestrians [2]. In order to understand how pedestrians will interact with AVs, it is essential to understand how pedestrians interact among themselves in interactive traffic scenarios. By understanding how pedestrians navigate and negotiate with fellow road users, insights can be gained that will enhance the development of safer and more efficient AVs. Investigating pedestrian-pedestrian interactions is therefore an essential step that will aid the integration of safe and effective AVs in our evolving urban landscapes.

To gain this understanding, different approaches have been employed. One of these approaches is trajectory prediction, where the future states of pedestrians, or dynamic agents in general, are predicted by a forecasting model based on their current and past states. The reasoning behind this is that if the predictions are accurate, collisions can be avoided by planning a path that is not in conflict with the anticipated movements of these dynamic agents. An alternative approach involves the use of modelling methods. These methods aim to create a model that describes the underlying dynamics and behaviours of agents. Unlike trajectory prediction methods, modelling methods try to capture the broader patterns and principles that guide human interactions.

AVs often rely on predictions of the future, which is commonly done with trajectory prediction. However, future human behaviour is inherently uncertain, both aleatoric i.e. by the randomness of external events, and epistemic i.e. by lack of knowledge from the observer [3]. High and lower-level decisions, related to behaviour levels [4], are cause for different modes. In this context, these modes refer to the multiple possible outcomes or paths of the future trajectory. An example would be a person walking towards an intersection and having options to go straight, left or right based on their desired destination. These high-level modes might be obvious, but there are also lower-level modes which might include slowing down or speeding up to pass before someone else, and making a tight or a wide turn to avoid a puddle. The complexity and situational dependence of these modes make it difficult to make accurate trajectory predictions for human road users.

To address the issue of uncertainty, numerous methods

aim to predict future agent trajectories probabilistically. These methods vary from using Gaussian Mixture Models (GMMs) [5, 6] to employing advanced generative networks. Generative networks like Generative Adversarial Networks (GANs) [7, 8], Conditional Variational Autoencoders (CVAEs) [9, 10], and Variational Recurrent Neural Networks (VRNNs) [11, 12] are particularly interesting as they can understand complex patterns without needing to know the number of expected modes. They offer the advantage of flexibility and the ability to handle various complexities within trajectory prediction tasks, as opposed to GMM-based methods. However, despite these cutting-edge techniques showing good prediction accuracy, they tend to suffer from the problem of mode collapse. Here the generator fails to capture the full diversity of the data distribution and produces limited or repetitive outputs, which is especially a problem in the GAN-based methods.

To tackle these challenges, a promising method is using Normalizing Flows (NFs) which are specifically designed to learn how data is distributed. A particularly promising method is TrajFlow [13], which is a probabilistic trajectory prediction approach focused on improving the accuracy of learned distributions and accurately predicting uncertain human behaviour in traffic scenarios.

Trajectory prediction methods like TrajFlow, need to be trained on trajectory data before they can predict trajectories. This training data has a crucial impact on the effectiveness of the model to create accurate predictions. There are a great number of real-world datasets that are used for trajectory training and evaluation, such as ETH/UCY [14, 15], SSD [16], nuScenes [17] and KITTI [18]. However, they all face a fundamental issue: only a single trajectory is observable out of all the potential future trajectories that could have been taken. As discussed previously, a single starting trajectory can result in multiple valid future trajectories influenced by high- and low-level modes. A solution for overcoming this issue is provided in the Forking Paths dataset [19]. By creating a simulated environment where annotators create multi-future trajectories, multimodality is established in the data.

Another reason why it is difficult to create an understanding of how human road users like pedestrians act in interactive scenarios is that in interactions, mostly described by modelling frameworks, both road users influence each other's behaviour by continuously responding to the actions of the other agent. These responses are guided by the communication between agents and determine priority and acceptance among road users [20]. Examples of this are a merging scenario where one vehicle needs to make way for the other, or two pedestrians walking on a collision path where both agents watch each other's movements in order to pass safely. Modelling these complex interactions proves to be a difficult task.

For the issue of describing interactions, multiple modelling frameworks have been introduced. Modelling human traffic behaviour has mostly focused on studying individuals, like how one vehicle follows other vehicles [21, 22], changes lanes [23, 24], or accepts gaps in traffic [25, 26]. These models often assume that one person responds to others, but those others do not respond back. For example, in car-following models, the following driver reacts to the leading vehicle, but the leading one does not react based on the follower. This one-way interaction assumption [27] helps to understand a single person's behaviour but is not ideal for scenarios involving mutual interactions, like intersections. Game theory still is commonly used to model these interactive agent systems [28, 29, 30]. However, it faces challenges as it assumes agents always act rationally, ignores the role of communication, and simplifies agents' actions.

To address these limitations, the Communication-Enabled-Interaction (CEI) framework [27] was proposed. This framework is built to model the entire two-way interaction between agents by considering the joint interactive system, guided by communication and the creation of beliefs about another person's intentions. It does not have the limitations of game-theoretic models and can comprehensively describe interactive behaviours.

Although complex trajectory prediction models like TrajFlow are great at handling the uncertainty and complexity inherent in pedestrian trajectories, they can fall short in capturing the nuances of mutual interactions through communicating intentions among pedestrians in interactive scenarios. On the other hand, while modelling frameworks like CEI do provide a sophisticated methodology to model interactions among agents, they lack a valid implementation for 2D scenarios, including a robust method that can create beliefs about complex and uncertain pedestrian trajectories. This difference highlights a significant gap in research. An approach is needed that joins the detailed understanding of how pedestrians interact using the CEI framework, and the advanced predictive abilities of TrajFlow. This approach will aim to improve the understanding of pedestrian behaviour in interactive situations by simulating pedestrian movements in interactive traffic scenarios. So, to bridge this gap the following research question is formulated:

What influence does the integration of TrajFlow's multimodal trajectory prediction in the Communication-Enabled-Interaction framework have on simulated pedestrian behaviour in interactive traffic scenarios?

To answer this question the CEI framework is combined with TrajFlow. The dynamics of the CEI model are expanded to suit 2D pedestrian simulations and will handle the overall interaction modelling, while the belief about other persons' actions is generated by TrajFlow multimodal trajectory prediction, making use of the

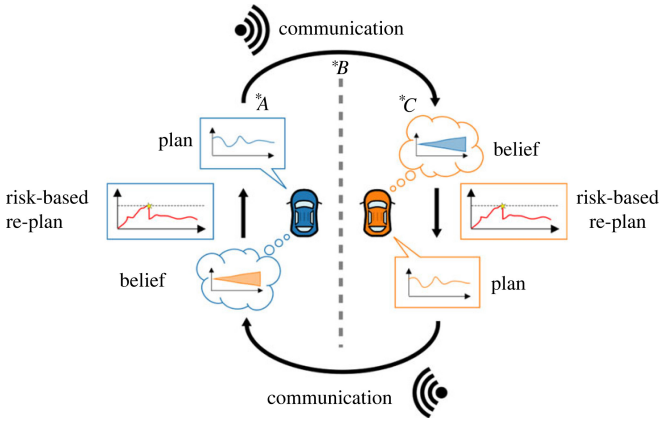


Figure 1: A schematic overview of the CEI framework [27]. Plans update based on risk-thresholds & risk estimates emerging from a belief of the other driver. Communication links one driver's plans to the belief of the other.

combined potential of these methodologies. TrajFlow is trained on the Forking Paths dataset, where its belief-based nature fits perfectly with TrajFlow's role in the CEI model. A pedestrian interaction simulation environment is designed, aiming to showcase its effectiveness.

## II. Background

### A. Communication-Enabled-Interaction Framework

The Communication-Enabled-Interaction (CEI) framework [27] as proposed by Siebinga et. al. forms a solid basis for modelling agent interaction behaviour. This framework considers that agents have plans and beliefs about others' actions, affecting their behaviour based on perceived risks. Unlike other models that assume individual, one-sided interactions between agents, the CEI framework captures the entire interaction between them. It provides a structured method for modelling human-human traffic interactions, and by considering the joint interactive system and integrating explicit communication, it does not have the limitations of traditional game-theoretic models, offering a more comprehensive understanding of interactive behaviours.

The CEI framework consists of four essential components describing the interaction between two agents: risk perception, a deterministic plan, communication, and probabilistic belief. A schematic overview is shown in figure 1.

*1) Risk-Based Re-plan:* The framework combines risk-based decision-making with the human preference for satisfactory solutions. Agents adjust their plans based on perceived risk levels, with low risk-threshold agents adapting early and high risk-threshold agents benefiting from others' risk reduction.

*2) Plan:* The plan defines an agent's intended actions for the near future covering part of the interaction. This plan, created by a risk-unrelated planning algorithm, prioritizes factors like desired speed and comfort, which is achieved by minimizing the following cost function:

$$c = \sum^{N} \left(v_n - v^d\right)^2 + \left(a_n^{\text{in}}\right)^2 \tag{1}$$

where $N$ is the time step, $v$ is the velocity, and $a^{\text{in}}$ is the input acceleration. Risk perception is regularly evaluated to determine plan suitability. If re-planning is needed, the risk-thresholds constrain the planning process.

*3) Communication:* The communication links one agent's plan to the other agent's belief. In the framework, an implicit communication model is used which observes the agent's position $p$ and velocity $v$.

*4) Belief:* Both agents are assumed to have probabilistic beliefs of the anticipated actions of the other agent in the near future. Bayesian updating is used where the previous belief point serves as the prior distribution, and the resulting posterior is used as the updated belief point. The likelihood is assumed to be a Gaussian distribution with a known standard deviation with position $p$, time $t$ and maximum comfortable acceleration $a_c$.

$$\mathcal{N}\left(\mu = \frac{p}{t}, \sigma^2 = \left(\frac{a_c t}{6}\right)^2\right) \tag{2}$$

The framework runs simulation steps with a time interval $\Delta t$. For each step, the agent updates its belief of the other agent and computes the perceived risk. If this risk is higher than the upper risk-threshold their plan is updated so that the perceived risk will be reduced. If the risk is below the lower risk-threshold and the last plan update was a longer time ago than the saturation time $\tau$, the plan is also updated to ensure the agent does not stay stuck in a plan that has a low risk, but does not move efficiently to their goal.

### B. TrajFlow

TrajFlow as proposed by Mészáros et. al. [13] is a probabilistic trajectory prediction method based on Normalizing Flows (NFs), that provides an analytical expression of the learned distribution. By using a Recurrent Neural Network Auto Encoder (RNN-AE) to encode trajectories into a lower-dimensional abstraction and NFs to learn distributions over these abstracted trajectory features, TrajFlow is able to capture complex, multi-modal distributions without the need to predefine the number of expected modes. This means it is applicable to a broad spectrum of scenarios and diverse datasets, offering flexibility and robustness in handling trajectory prediction tasks. It has shown predictive performance on par with or superior to state-of-the-art methods on the real-world ETH/UCY [14, 15] pedestrian dataset, supporting the previous statement. Furthermore, TrajFlow distinguishes itself by offering open-source code for its framework, enabling easy access, modification, and application of the model to specific domains and datasets. A schematic overview is given in figure 2.

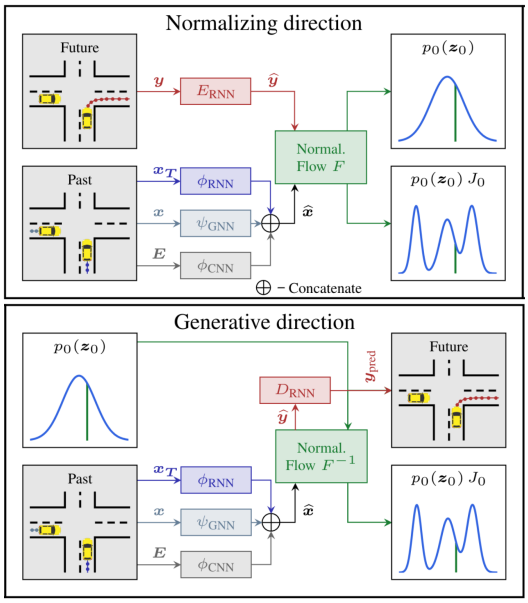


Figure 2: Schematic overview of the TrajFlow architecture [13]. In the normalizing direction trajectory information is encoded to learn the distribution with NFs, in the generative direction this distribution is used to generate new future trajectories.

### C. Forking Paths

The Forking Paths dataset [19] is constructed by taking real-world pedestrian datasets and recreating them in the 3D CARLA simulator [31], allowing multiple annotators to control pedestrians. The result is a truly multimodal dataset containing multiple future trajectories branching from the exact same starting trajectory, containing multiple variations in high-level human behaviour, which is impossible to simulate automatically. The set contains 750 sequences from 10 annotators across 7 scenes and provides diverse annotations for identical past trajectories. Each sequence covers about 15 seconds, with an average of 5.9 future trajectories per agent. See figure 3 for an example.

## III. Methods

In order to create a framework capable of showing realistic pedestrian behaviour in interactive pedestrian scenarios, CEI, TrajFlow, and Forking Paths were utilized. The CEI framework was modified and formed the base environment in which everything was combined, as this framework provides a structured method for modelling two-sided interactions with its four modules: belief, plan, risk and communication. TrajFlow was used to formulate the beliefs of the pedestrians, as it can capture the complex, multi-modal distributions present in pedestrian movements and use that to predict trajectories accompanied by a likelihood. The Forking Paths dataset aligns seamlessly with the study's objectives as it not only

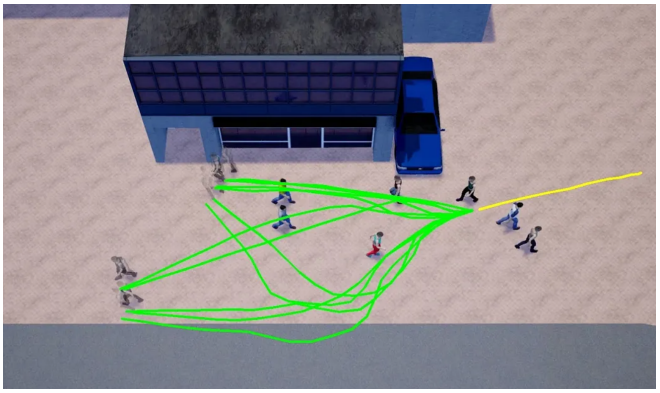

Figure 3: Example from the Forking Paths dataset [19]. The yellow line denotes the past trajectory of a pedestrian and the green lines are continuations of the annotators in the simulation.

contains multimodal pedestrian data but also consists of trajectories that represent the beliefs of the annotators, making them highly suitable as the foundation for constructing beliefs within the CEI framework. Therefore TrajFlow was trained on the Forking Paths dataset.

### A. Implementation

TrajFlow was trained on the Forking Paths dataset, where 8 timesteps ($\Delta t_{tf} = 0.4s$) leading up to the split in multiple futures were taken as past trajectories and 12 timesteps after the split were taken as future trajectories. The past trajectories of all pedestrians were encoded and concatenated to form the conditional argument of the NF. The future trajectories were encoded separately and were what the NF aimed to learn the distribution of. After this distribution was learned, for prediction, TrajFlow encodes the past trajectory information of the pedestrians in the scene, and with the learned distribution was able to output 100 possible future trajectories together with a likelihood score. The Forking Paths data that was used for the training of TrajFlow consists of all the scenes where only pedestrians interact with each other. The positions from a top-down view were extracted and scaled to be represented in meters.

### B. Modification

Several changes have been made in order to utilise the CEI framework for a two-dimensional pedestrian scenario simulation. These changes are in the overall dynamics of the system, as well as planning, belief construction, and risk evaluation.

*1) Overall dynamics:* In the paper where the CEI framework was proposed, a case study was performed that tests the effectiveness in a simple merging scenario. This scenario consisted of two cars modelled as point masses moving along a predefined track where the accelerations, velocities, and positions were all expressed in one dimension.

$$
\begin{aligned}
a_x^{net}(v_x) &= a^{in}(v_x) - a^r(v_x) \\
v_x^{t+1} &= v_x^t + a_x^t * t \\
p_x^{t+1} &= p_x^t + v_x^t * t
\end{aligned} \tag{3}
$$

$a^{net}$ is the net acceleration, $a^{in}$ is the applied input acceleration, $a^r$ is the negative acceleration due to friction forces, $v$ is velocity, and $p$ is position. Since the goal of this research is to examine the system for pedestrians in two dimensions, the point mass dynamics were adjusted so that accelerations, velocities, and positions are also expressed in 2D. The frictional forces have been neglected, as their effect on the net accelerations of humans walking at a low velocity is minimal.

$$
\begin{aligned}
a_{x,y}^{net}(v_x, v_y) &= a^{in}(v_x, v_y) \\
v_{x,y}^{t+1} &= v_{x,y}^t + a_{x,y}^t * t \\
p_{x,y}^{t+1} &= p_{x,y}^t + v_{x,y}^t * t
\end{aligned} \tag{4}
$$

*2) Planning:* The planning is performed in acceleration space. Differing from the original model, each pedestrian has a maximum and minimum acceleration $a_{max}$ and $a_{min}$ in both $x$ and $y$ directions. In the planning step, an optimizer searches actions between those minima and maxima for the prediction time horizon $T$. Using the 2D point mass dynamics from equation 4 it can infer the velocities and positions for all $T$.

Since the goal is not anymore for a car to follow a predefined path comfortably, but rather for pedestrians to reach their end position, the cost function from equation 1 was altered.

$$
\begin{aligned}
c &= c_p + c_v \\
c_p &= \sum_{n=1}^{T} \left( \|p_n - p_{goal}\|_2 \right)^2 \\
c_v &= \sum_{n=1}^{T} \left( \begin{cases} 0 & \text{if } v_{diff} \leq 0 \\ v_{diff}^2 & \text{if } v_{diff} > 0 \end{cases} \right) \\
v_{diff} &= \|v_n\|_2 - v_{pref}
\end{aligned} \tag{5}
$$

This cost function $c$ consists of a positional part $c_p$ and a velocity part $c_v$. The positional cost is calculated by taking the L2 norm (calculating the absolute distance) of the difference between all future positions $p_n$ and the goal position $p_{goal}$ of the pedestrian, and taking the sum. Positions closer to the goal yield a lower positional cost and this gives the pedestrian incentive to move towards the goal position. The velocity cost is calculated by taking the L2 norm of all the velocities $v_n$ for time $T$, subtracting the preferred velocity $v_{pref}$, setting the cost to zero when it is below zero and squaring the remaining values after which the sum is taken. This essentially means that velocities with a norm higher than the preferred (average walking) velocity are penalized exponentially while slowing down to prevent a collision or when the destination is reached will never be penalised. Within the bounds of risk, this optimizer searches for the optimal acceleration actions to minimize the cost function.

*3) Belief construction:* The belief construction in the original framework was executed by observing the other agent's 1D position and velocity on the predefined track, extrapolating it over time assuming constant velocity, resulting in a Gaussian distribution with a constant and known standard deviation. For the 2D pedestrian implementation, trajectories become more uncertain and complex due to multimodality, and a more sophisticated method of creating beliefs is necessary. Therefore belief construction is handled differently with TrajFlow. Every time the belief is updated, TrajFlow receives the last eight positions $p_{[t_{tf}-7,t_{tf}]}$ with an interval of $\Delta t_{tf}$ of the other agents' history and returns 100 future trajectories of twelve position points $p_{[t_{tf}+1,t_{tf}+12]}$ accompanied by log probabilities for each trajectory. These trajectories consist of possible future positions for the other pedestrian and form the belief that is used in the risk evaluation. To find out how TrajFlow's belief construction performs, another simpler method for creating beliefs was used for comparison. This belief construction is based on the assumption of constant velocity and acceleration, similar to the 1D simple merging scenario, infused with Gaussian noise.

$$
\begin{aligned}
p_{new} &= p_{old} + v * dt + 0.5 * a^2 * dt \\
\text{noise} &\sim \mathcal{N}(0, \frac{1}{20}) \\
p_{new} &= p_{new}\text{noise} * \left(1 + \frac{\text{belief\_index}}{\text{belief\_length}}\right) \\
p_{old} &= p_{new}
\end{aligned} \tag{6}
$$

First, $p_{new}$ is calculated with kinematics, and then a Gaussian noise is sampled, which is added to $p_{new}$ proportional to the distance of the point. Lastly, $p_{old}$ is updated and this is repeated until there are 100 trajectories with 12 timesteps each, conforming to the shape of the TrajFlow belief. This belief based on constant velocity and noise will be referred to as the naïve belief from this point forward.

*4) Risk evaluation:* Since the risk evaluation is based on both the plan and belief, and both elements have been modified, a new method to evaluate the risk had to be used.

During the risk evaluation, the pedestrian's position plan is compared with the beliefs. More precisely, the planned positions of the pedestrian in the time horizon $T$ are filtered and matched for the timestamps $[t_{tf} + 1, t_{tf}+12]$ from the belief and compared with the 100 belief positions from TrajFlow. A distance array $d$ is formed containing all the distances between the twelve positions within the planned trajectory and the 100 belief positions at the corresponding timesteps.

$$
d = \|p^{plan}_{[t_{tf}+1,t_{tf}+12]} - p^{belief}_{[t_{tf}+1,t_{tf}+12]}\|_2 \tag{7}
$$

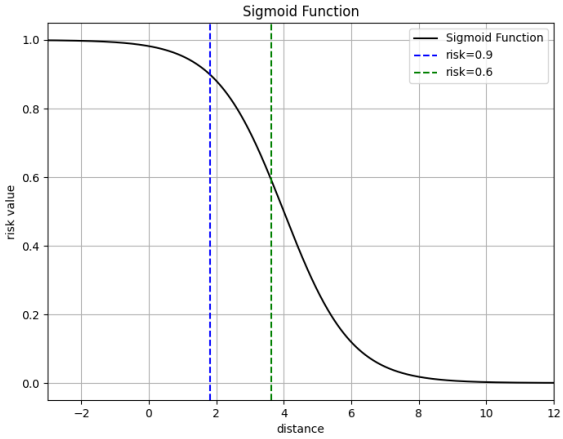


Figure 4: Sigmoid function for computation of risk values.

After the absolute distances between the plan and belief positions are computed, they are passed to a Sigmoid function shown in figure 4.

$$\sigma(d) = \frac{1}{1 + e^{d-4}} \tag{8}$$

This function returns a risk of collision between 0 and 1, scaled so that the risk goes to 0 as the distance is great, then increases with an S-shape going to 1 when the distance nears 0. The risks for the 100 belief trajectories and position plan at corresponding timesteps are multiplied by the normalised probabilities $P$ for each belief trajectory and summed, resulting in a risk array $r_{array}$ with a single risk for each of the 12 position plan points. The maximum risk is taken as the risk $r_{plan}$ of the current plan, which will trigger a re-plan if necessary based on the risk-thresholds.

$$\begin{aligned} r_{array} &= \sum_{n=1}^{100} (\sigma(d_n) * P) \\ r_{plan} &= \max \{r_{array}\} \end{aligned} \tag{9}$$

Evaluating a plan on the maximum perceived risk of that plan is based on the intuition that a person will not continue to carry out a plan for which they feel the risk is too high at some point and is performed similarly in the 1D simple merging scenario. An alteration was made to the framework so that if the last re-plan was performed because the upper risk-threshold was exceeded and afterwards the risk has dropped under the lower risk-threshold, it does not wait for $\tau$ seconds before re-planning but instead performs a single re-plan directly before going back using the saturation time normally. This ensures that the pedestrians will not keep following a non-optimal path after the risk of collision has disappeared.

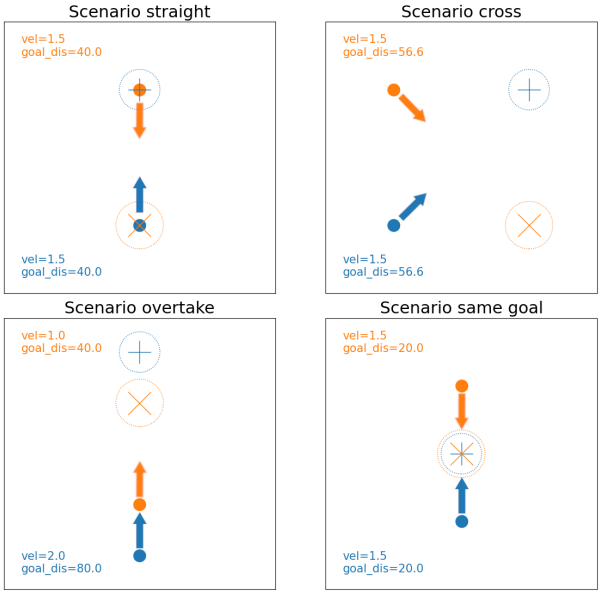


Figure 5: Overview of the scenarios: The pedestrians are shown as dots and their goal positions as dotted circles with a cross in the middle. All pedestrians have an initial heading in the direction of their goal positions.

### C. Experimental Setup

*1) Simulations scenarios:* A simulation environment was created to emulate pedestrian interactions and assess the effectiveness of the combined model. This pedestrian simulation environment is a 2D plane where two pedestrians can move freely, considering their dynamics. Four different walking scenarios were created within this simulation environment. These scenarios were selected because they all contain conflicting trajectories and therefore two-way interaction is necessary to resolve these conflicts. Each scenario mirrors real-life situations, like avoiding collisions or heading toward the same goal. These scenarios let us test how TrajFlow and the CEI framework model simulate pedestrian behaviour and assess their effectiveness in modelling these interactions. A visualisation of these scenarios is shown in figure 5.

*a) Walking straight scenario:* In this scenario, both pedestrians are placed at a certain distance from each other. The pedestrians have the same absolute velocity and are walking towards each other. Their goal position is the starting position of the other pedestrian. In this experiment, the pedestrians are forced to move out of each other's way in order to reach their destination without colliding.

*b) Crossing paths scenario:* In this scenario, both pedestrians are placed on the left in the simulation. They have the same absolute velocity and a heading so that they will cross paths in the middle of their trajectory to their

goal position with an angle of 90°. In this experiment, the pedestrians will need to avoid collision by either moving out of the way or changing their velocity to pass before or after the other pedestrian.

*c) Overtake scenario:* In this scenario, one pedestrian is trying to overtake the other. The first pedestrian has an absolute velocity twice as large as that of the other. The goal position is in the same direction but twice as far away. Trying to maintain their preferred velocity, one has to overtake the other.

*d) Same goal scenario:* In this scenario, both pedestrians are trying to reach the same goal position, have the same absolute velocity and are walking towards each other. This will result in conflicting behaviour, where the risk-bounds play a key role in which pedestrian will yield first.

*2) Parameters & data collection:* The simulations were performed in the CEI framework with a timestep of $\Delta t = 0.05s$ and a saturation time $\tau = 1.0s$. Different risk-thresholds were selected for the simulations using TrajFlow belief and the ones using the naïve belief. Non-identical belief systems lead to non-identical perceived risks, which inevitably lead to differences in behaviour. Simply comparing behaviours then is unfair, which is why the risk-thresholds were selected to have a similar perceived risk in identical situations. For finding suitable risk-thresholds a simple test simulation was performed in which the pedestrians did not update their plans. The perceived risks of the pedestrians for both belief constructions were plotted over time as can be seen in figure 6. This provides a point at which the upper risk-thresholds can be selected so that both thresholds will be exceeded at the same time. For the TrajFlow-based belief, the lower and upper risk-thresholds were set at 0.6 and 0.8 respectively and for the naïve belief, the lower and upper risk-thresholds were set at 0.7 and 0.9 respectively. To learn how the framework behaved when two pedestrians did not have the same risk-thresholds, simulations were run where one of the pedestrians had its risk-thresholds reduced by 0.3 for both the upper and lower threshold. Data collection was performed by running the simulations described in section III-C1. For each of the four scenarios, with TrajFlow belief or naïve belief, with equal or different risk-thresholds, 10 simulations were executed resulting in a total of 160 simulations. A full overview is shown in Appendix A. Although the CEI model itself is deterministic, the belief constructions are not. It is therefore important to run multiple simulations so that outliers will not negatively impact the overall analysis. For all simulations, data is recorded at each timestep.

*3) Metrics:* In order to create quantifiable results, metrics were constructed that helped describe the simulation results. These metrics were used to see whether the framework as a whole behaves as expected and compare TrajFlow's and the naïve beliefs.

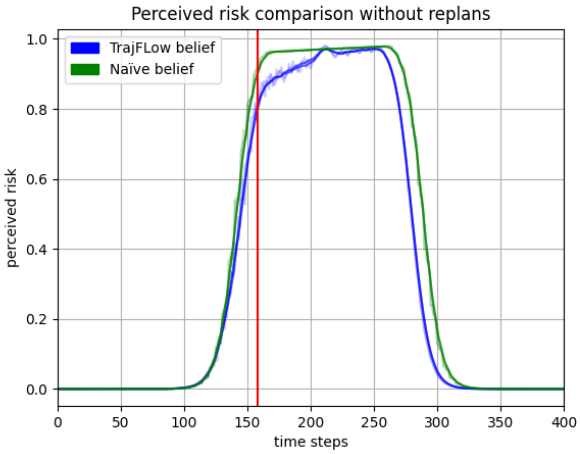


Figure 6: Perceived risk of pedestrians over time when no re-plans are performed in the straight scenario.

*a) Maximum deviation:* For each trajectory, the maximum deviation from the from the shortest path from start to goal was calculated as follows:

$$\begin{gathered}\text{For } i = 0 \text{ to } n \\ \text{dev}_i = \frac{|(x_2 - x_1)(y_1 - y_i) - (x_1 - x_i)(y_2 - y_1)|}{\sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}} \\ \text{maximum deviation} = \max\{\text{dev}_i\}\end{gathered} \tag{10}$$

$\text{dev}_i$ is an array containing all the deviations, where $(x_1, y_1)$ are the start coordinates, $(x_2, y_2)$ are the goal coordinates, and $(x_i, y_i)$ are the $i^{\text{th}}$ coordinates from the path trajectory. This metric shows how close the simulated pedestrians will stay to their preferred path. The expectation is that pedestrians with a higher risk-threshold will stay closer to their preferred path compared to pedestrians with a lower risk-threshold.

*b) Moment of deviation:* For each pedestrian, the moment of deviation is also computed. This is the first index where the deviation from the shortest path to the goal is greater than $0.1m$.

$$\text{moment of deviation} = \min\{i \mid [\text{dev}_i > 0.1]\} \tag{11}$$

This metric is especially interesting when the pedestrians do not have the same risk-threshold. If the pedestrian with the lower risk-threshold will yield earlier than the one with a higher risk-threshold, the framework is working as intended.

## IV. Results

This chapter shows the results that follow from the experiments. For each scenario, the expectations are given and the observations are discussed. Each scenario contains four combinations, TrajFlow or naïve belief construction with equal or different risk-thresholds, containing ten

simulations each. In all simulations where the risk-thresholds are different, the orange pedestrian has the lower risk-threshold values. Visualising all this data from 160 simulations here is not feasible so examples are used to visualise the results. These results are handpicked and contain favourable behaviour. Unfavourable behaviour and other inconsistencies are discussed. For the complete visualisations, please refer to Appendix B.

### *A. Straight Scenario*

*1) Expectations:* It is expected that both pedestrians will walk toward each other before yielding and continuing to their destination. Since this scenario is symmetric, the expectation for pedestrians with equal risk-thresholds is that both pedestrians yield an equal amount for each other on average. Note that the belief construction is stochastic, so variations are expected. When the pedestrians have different risk-thresholds it is expected that the low risk-threshold pedestrian yields earlier and more on average. In both cases, the magnitude and direction of yielding will depend on the shape of the constructed belief.

*2) Observations:* Visualisations of the examples are shown in figure 7. Let's begin by examining the figures 7a and 7e where the risk-thresholds were the same for both pedestrians and the belief was created with TrajFlow. Conforming to the expectation, the pedestrians walked towards each other, before yielding a comparable amount, after which they followed the fastest path towards their goals. This behaviour was reflected in the rest of the simulations (Appendix B.1), where in most cases both pedestrians yielded a comparable amount, with some cases where one yielded more than the other as a result of the stochasticity.

Next, we'll examine the figures 7b and 7f, where the risk-thresholds were also the same for both pedestrians, but the naïve belief construction was used. The trajectories looked very similar to the first ones, where both pedestrians walked towards each other before yielding a comparable amount and continuing to their goals. A noticeable difference was that the pedestrians yielded later and more abruptly. When looking at the other nine trajectories (Appendix B.1), some cases with unnatural behaviour emerged. This includes cases where pedestrians yielded more than five meters, walked backwards or even walked in a small circle.

For the third configuration, figure 7c and 7g, the risk-thresholds were different for both pedestrians and the belief was created with TrajFlow. The behaviour was as expected, where the high risk-threshold pedestrian stayed closer to its preferred path. Looking at the other nine trajectories, this behaviour showed in most cases (Appendix B.1). The pedestrian with the high risk-threshold had a significantly lower maximum deviation and the pedestrian with the low risk-threshold deviated

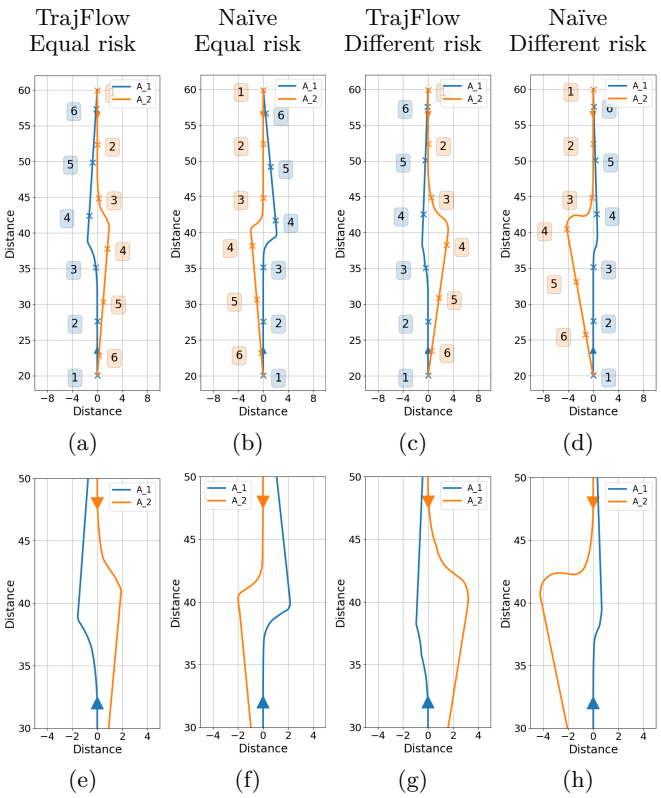


Figure 7: Straight scenario trajectory examples. 7a uses TrajFlow belief construction and has equal risk-thresholds, 7b uses the naïve belief and has equal risk-thresholds, 7c uses TrajFlow belief and has different risk-thresholds, 7d uses the naïve belief and has different risk-thresholds. When the risk-thresholds are different, the orange pedestrian has lower risk-thresholds. Underneath in 7e, 7f, 7g, 7h the same scenarios are shown zoomed in on the interaction with equal x-axis and y-axis proportions.

significantly earlier from their preferred path, as shown in figure 8a and 8b.

For the last configuration, figure 7d and 7h, the risk-thresholds were different for both pedestrians and the naïve belief construction was used. The trajectories had a similar shape to the ones with TrajFlow belief, where the high risk pedestrian stayed closer to its preferred path. Looking at the other nine trajectories (Appendix B.1), it was observed that the high risk pedestrian stayed closer to its preferred path as can be seen in figure 8c. The low risk pedestrian however did not deviate significantly earlier, and actually moved backwards in all of the cases.

*3) Beliefs:* The observed behaviour described above can be attributed to the difference in belief construction. Examples of the data-driven TrajFlow belief and the naïve belief are shown in figure 9. The belief created by TrajFlow, where beliefs with a higher likelihood are

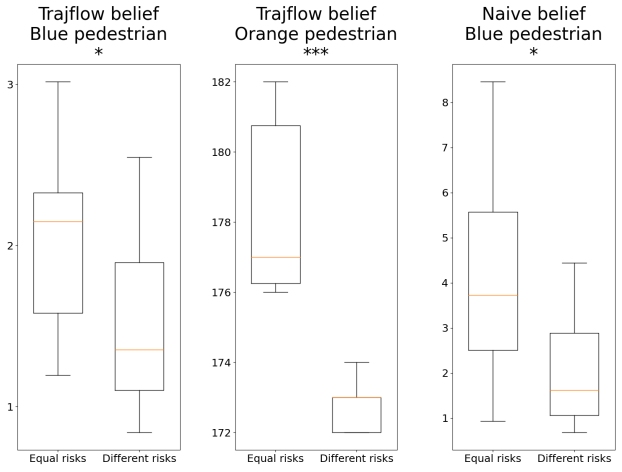


(a) (b) (c)

Figure 8: Boxplots comparing metrics of pedestrians between equal and different risk-thresholds in the straight scenario.

plotted in darker green, clearly shows the expectation that the pedestrian will try to avoid and move away from the straight path to one side or the other. The naïve belief construction does not show a preference to move in any other direction than continuing straight with Gaussian-infused noise. The naïve belief also has a narrower spread compared to the data-driven belief. A consequence of this is that in the simulations containing the naïve belief, the perceived risk will stay low for longer before suddenly increasing quite sharply, explaining the late and more abrupt deviations of trajectories.

*4) Overall influence TrajFlow:* Visually comparing the trajectories of the framework with TrajFlow belief and naïve belief in this scenario demonstrates the benefits of using TrajFlow. The pedestrians show plausible interaction behaviour, where pedestrians with equal risk-thresholds both yielded a comparable amount in order to safely pass each other. When the risk-thresholds were different the pedestrian accepting low risk yielded earlier and the pedestrian with the higher risk-threshold deviated less from its preferred path. The framework with the naïve belief lacked this behaviour, where with the same risk-thresholds the pedestrians sometimes passed each other on the same side of their preferred path, made more sudden turns and sometimes even moved backwards or in a circle to avoid collision. When the risk-thresholds were different the pedestrian accepting higher risk still had to yield substantially. The low risk-threshold pedestrian did not deviate earlier and moved backwards in all ten simulations.

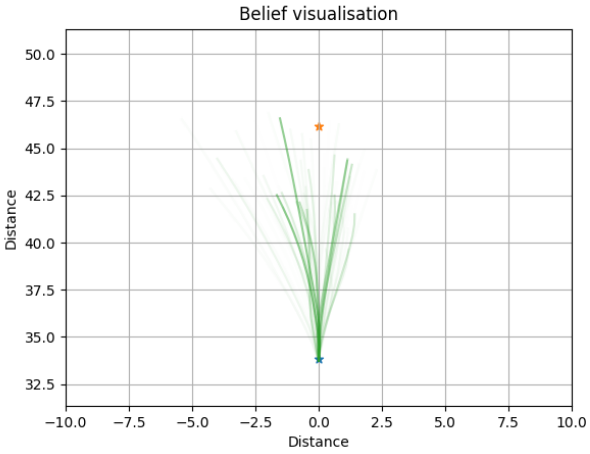


(a) TrajFlow belief

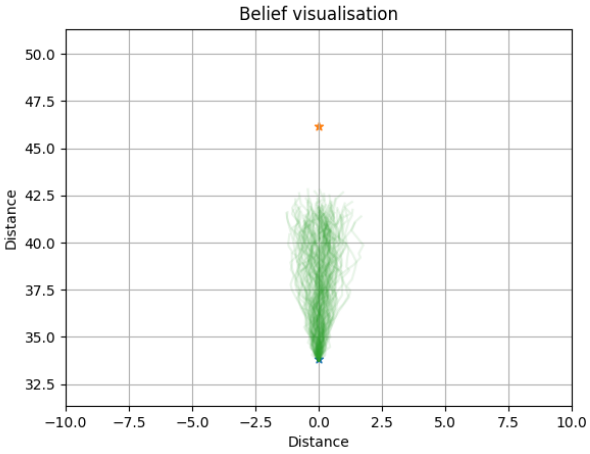


(b) Naïve belief

Figure 9: Visualisation of TrajFlow belief 9a and naïve belief 9b in the straight scenario.

## B. Cross Scenario

*1) Expectations:* The expectation for this scenario is that both pedestrians will walk toward the centre where they will meet. The scenario is symmetric so it is expected that pedestrians with equal risk-thresholds will deviate a comparable amount on average in order to safely pass each other. Variations are expected since the belief construction is stochastic. When the pedestrians have different risk-thresholds the expectation is that the low risk-threshold pedestrian will yield earlier and more on average. In both cases, the magnitude and direction of yielding will depend on the shape of the constructed belief.

*2) Observations:* Visualisations of the examples are shown in figure 10. Let's begin by examining the figures 10a and 10e where the risk-thresholds were the same for both pedestrians and the belief was created with TrajFlow.

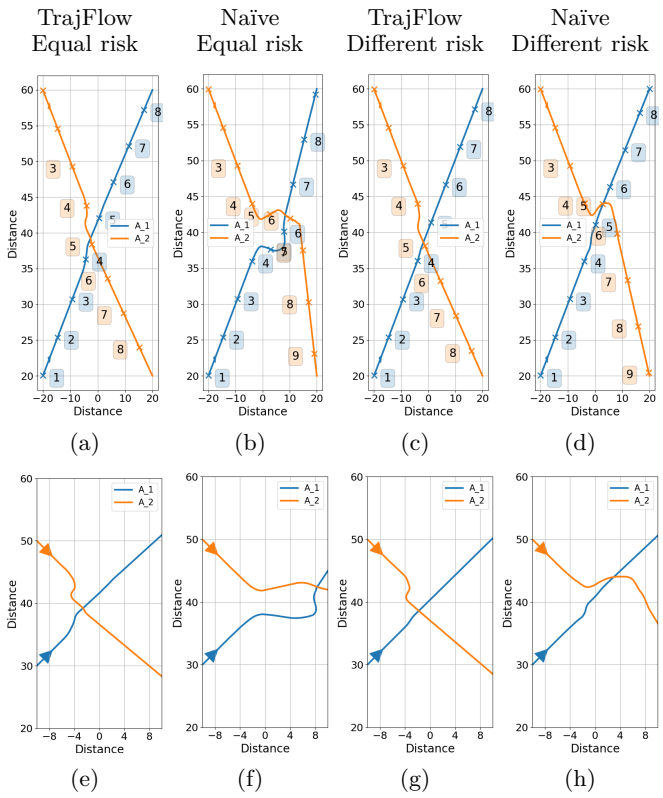


Figure 10: Cross scenario trajectory examples. 10a uses TrajFlow belief construction and has equal risk-thresholds, 10b uses the naïve belief and has equal risk-thresholds, 10c uses TrajFlow belief and has different risk-thresholds, 10d uses the naïve belief and has different risk-thresholds. When the risk-thresholds are different, the orange pedestrian has lower risk-thresholds. Underneath in 10e, 10f, 10g, 10h the same scenarios are shown zoomed in on the interaction with equal x-axis and y-axis proportions.

Conforming to the expectation, the pedestrians walked towards the centre, before yielding a comparable amount trying to pass behind the other, after which they followed the fastest path towards their goals. This behaviour was reflected in the rest of the simulations (Appendix B.2), where in most cases the pedestrians yielded to pass behind the other, sometimes letting the bottom pedestrian pass first, and sometimes the top one.

Next, we'll examine the figures 10b and 10f, where the risk-thresholds were also the same for both pedestrians, but the naïve belief construction was used. The trajectories also conformed to expectations in that the pedestrians walked towards the centre before yielding to resolve conflict, and continuing the fastest path towards their goals. A noticeable difference is that the pedestrians were trying to pass in front of each other. The same behaviour was shown in the rest of the simulations (Appendix B.2), where the paths often crossed in the right half of the simulation environment.

For the third configuration, figure 10c and 10g, the risk-thresholds were different for both pedestrians and the belief was created with TrajFlow. In this example the behaviour was as expected, where the most effort to avoid a collision was done by the low risk-threshold pedestrian, who passed behind the high risk-threshold pedestrian. Looking at the other trajectories (Appendix B.2), the behaviour of the high risk-threshold pedestrian staying closer to its preferred path was not dominant in all of the simulations, which is also reflected by the maximum deviation metric, which shows no significant change compared to same risk-threshold pedestrians. The moment of deviation was significantly later for the high risk-threshold pedestrian, which does conform with expectations and is shown in figure 11a.

For the last configuration, figure 10d and 10h, the risk-thresholds were different for both pedestrians and the naïve belief construction was used. The trajectories were as expected, where the most effort to avoid a collision was done by the low risk-threshold pedestrian, who passed in front of the other pedestrian on the right side of the simulation environment. Looking at the other nine trajectories (Appendix B.2), the behaviour of the high risk-threshold pedestrian staying closer to its preferred path was not dominant in all of the simulations. This is reflected by the maximum deviation metric, which shows no significant change compared to same risk-threshold pedestrians. There were also two cases in which the low risk-threshold pedestrian walked in a loop, in one case even causing the optimiser to fail and run off the page. The moment of deviation was significantly earlier for the low risk-threshold pedestrian, which does conform with expectations and is shown in figure 11b.

*3) Beliefs:* The observed difference in behaviour described above can be attributed to the difference in belief construction. Examples of the data-driven TrajFlow belief and the naïve belief are shown in figure 12. The belief created by TrajFlow, where beliefs with a higher likelihood are plotted in darker green, has a wider spread compared to the naïve belief. As a consequence, if the orange pedestrian tries to reduce its risk with the TrajFlow belief by moving a bit more to the right, the risk reduction is not that large, moving more down however, will decrease the risk more significantly. It might seem like the pedestrian goes right into the path of the belief, but note that the perceived risk is computed by comparing the current plan at a future timestep against the beliefs at that same future timestep. With the naïve belief, the perceived risk is low for a while and then suddenly increases at the end of its plan. Moving a bit more to the right will decrease that risk and make the plan safe again.

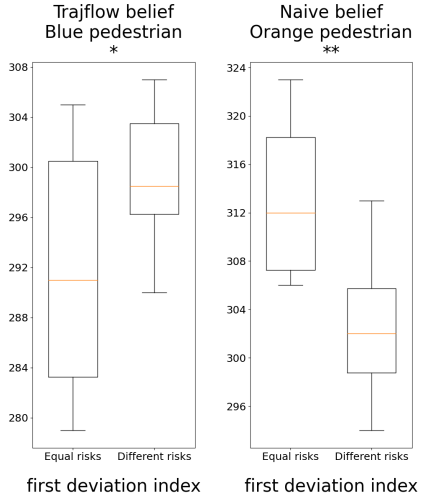


(a) (b)

Figure 11: Boxplots comparing metrics of pedestrians between equal and different risk-thresholds in the cross scenario.

*4) Overall influence TrajFlow:* Visually comparing the trajectories of the framework with TrajFlow and naïve beliefs in this scenario demonstrates a slight preference for using TrajFlow over the naïve belief implementation, although both do not show perfect behaviour. The simulations with TrajFlow show more plausible interaction behaviour. When the risk-thresholds were equal the pedestrians did not diverge far from their preferred path. After moving towards each other, one let the other pass before they continued to their goal. There were however cases in which one of the pedestrians moved backwards to resolve the conflict. When the risk-thresholds were different the pedestrian accepting high risk deviated later, but there was no significant difference in the amount of deviation and the expected behaviour is not shown in all the simulations. The framework with naïve beliefs performed a bit worse, where with the same risk-thresholds in 7 cases the pedestrians were unable to resolve the conflict quickly, resulting in the last part of the trajectory being almost straight up or down towards their goal. When the risk-thresholds were different the pedestrian with the lower risk-threshold yielded earlier, but there was no significant change in the maximum deviation. Furthermore, there are two cases in which the low risk-threshold pedestrian walked in a circle.

### C. Overtake Scenario

*1) Expectations:* The expectation for this scenario is that both pedestrians will walk upward towards their goal until a conflict has to be resolved. This scenario is not symmetric since the blue pedestrian has twice the speed and distance to cover to the goal. It is expected

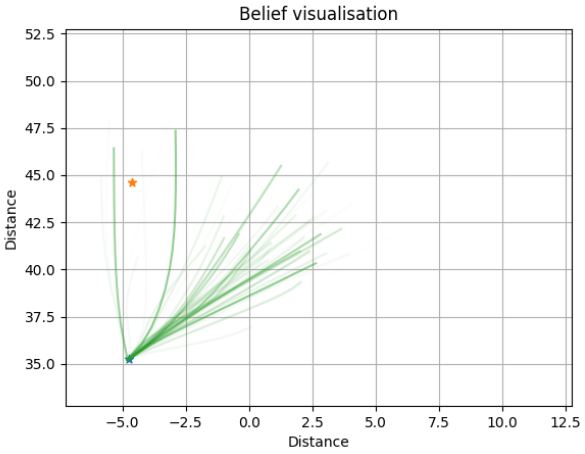


(a) TrajFlow belief

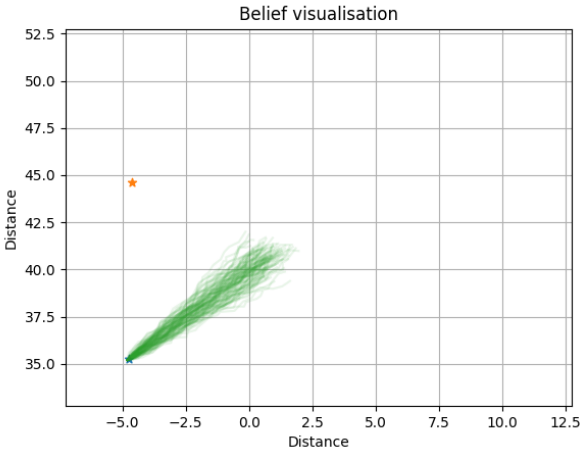


(b) Naïve belief

Figure 12: Visualisation of TrajFlow belief 12a and naïve belief 12b in the cross scenario.

that in the scenario where both pedestrians have equal risk-thresholds, the overtaking blue pedestrian will do (most of) the work to yield and avoid the other orange pedestrian. When the risk-threshold of the orange pedestrian is lower than that of the blue one, it is expected that the orange pedestrian will also move out of the way to make space for the high risk-threshold pedestrian. In both cases, the magnitude and direction of yielding will depend on the shape of the constructed belief.

*2) Observations:* Visualisations of the examples are shown in figure 13. Let's begin by examining the figures 13a and 13e where the risk-thresholds were the same for both pedestrians and the belief was created with TrajFlow. The expected behaviour was shown, where both pedestrians were walking straight towards their goal until the conflict needed to be resolved and the

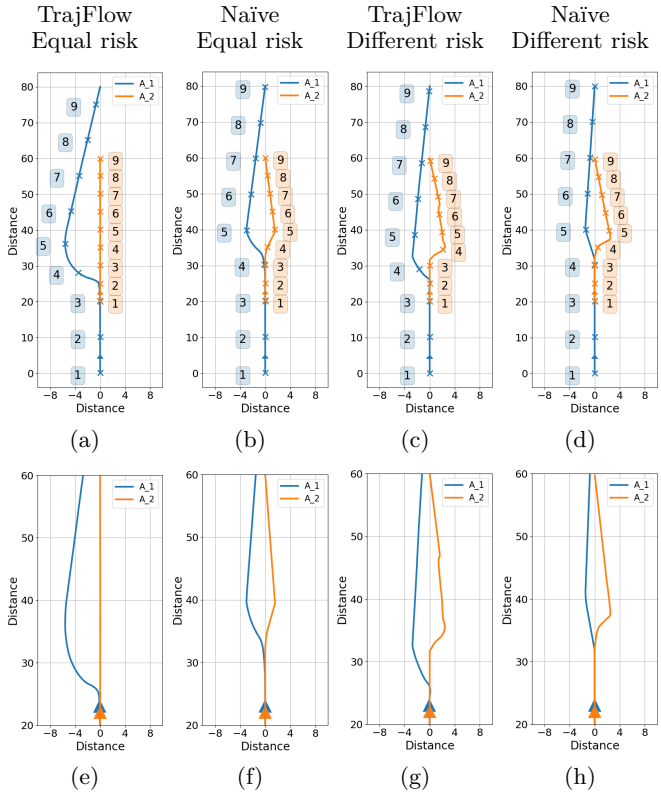


Figure 13: Overtake scenario trajectory examples. 13a uses TrajFlow belief construction and has equal risk-thresholds, 13b uses the naïve belief and has equal risk-thresholds, 13c uses TrajFlow belief and has different risk-thresholds, 13d uses the naïve belief and has different risk-thresholds. When the risk-thresholds are different, the orange pedestrian has lower risk-thresholds. Underneath in 13e, 13f, 13g, 13h the same scenarios are shown zoomed in on the interaction with equal x-axis and y-axis proportions.

blue pedestrian deviated from its path to pass the orange pedestrian while the orange pedestrian stuck to its preferred path. This behaviour was reflected in the rest of the simulations (Appendix B.3), where the overtaking blue pedestrian deviated from its path to perform the overtaking manoeuvre and the orange pedestrian stayed close to its preferred path.

Next, we'll examine the figures 13b and 13f, where the risk-thresholds were also the same for both pedestrians, but the naïve belief construction was used. Both pedestrians started walking towards their goal just like with TrajFlow belief. When the conflict started, both the overtaking blue pedestrian and the orange pedestrian yielded from their path, and looking closely at the timestamps shows that the orange pedestrian yielded even a bit earlier. In all the other nine cases (Appendix B.3), the orange pedestrian yielded more than the overtaking blue pedestrian and in four cases the blue pedestrian barely yielded at all. In all cases the orange pedestrian yielded first.

For the third configuration, figure 13c and 13g, the risk-thresholds were different for both pedestrians and the belief was created with TrajFlow. In this example the behaviour was as expected, where instead of just the overtaking blue pedestrian, now also the orange pedestrian made some effort to yield from its preferred trajectory. Looking at the other simulations (Appendix B.3), in all the cases the overtaking blue pedestrian made the most effort to yield and overtake, and in six cases the orange pedestrian also made an effort to yield. In two of those cases however, both pedestrians yielded in the same direction. The expected behaviour is reflected by metrics where the overtaking blue pedestrian yielded significantly later and the orange pedestrian yielded significantly more, compared to the situation where the risk-thresholds were equal, as is shown in figure 14a and 14b.

For the last configuration, figure 13d and 13h, the risk-thresholds were different for both pedestrians and the naïve belief construction was used. In the example the blue overtaking pedestrian yielded less than the orange pedestrian, which also yielded earlier from its preferred path. When taking a look at the full set of simulations (Appendix B.3), it was shown that the orange pedestrian yielded earlier and more in all cases, and in only three cases the blue overtaking pedestrian made any real effort to yield from its preferred path. The maximum deviation metric shows a significant increase for the orange pedestrian when the risk-thresholds were different as can be seen in figure 14c.

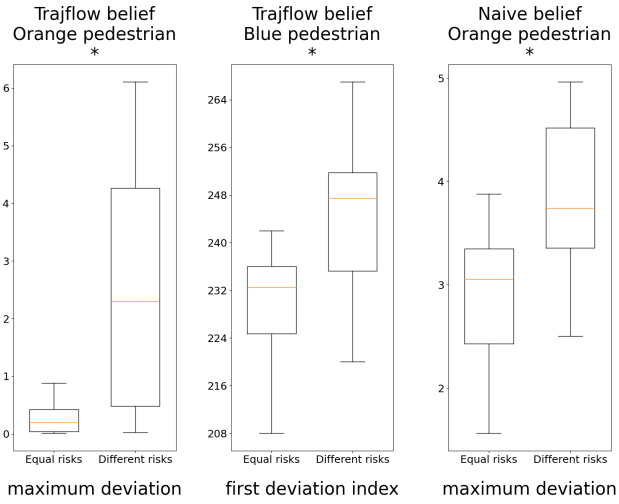


Figure 14: Boxplots comparing metrics of pedestrians between equal and different risk-thresholds in the overtake scenario.

*3) Beliefs:* The observed difference in behaviour described above can be attributed to the difference in belief construction. Examples of the data-driven TrajFlow belief and the naïve belief are shown in figure 15. The belief created by TrajFlow, where beliefs with a higher likelihood are plotted in darker green, clearly shows the expectation that the blue pedestrian will try to avoid and move away from the straight path to overtake on one side or the other. This results in the perceived risk of the orange pedestrian staying low, so it can continue on its preferred path while the blue pedestrian manoeuvres around it. The naïve belief construction has a narrower spread and does not show an expectation that the blue pedestrian will move away. As a consequence, the perceived risk of the orange pedestrian will suddenly increase quite steeply, triggering re-plans and yielding, allowing the blue overtaking pedestrian to stay closer to its preferred trajectory.

*4) Overall influence TrajFlow:* Visually comparing the trajectories of the framework in this scenario demonstrates the benefit of using TrajFlow. The pedestrians show plausible interaction behaviour, where in scenarios with equal risk-threshold pedestrians, the person coming from behind performing the overtaking manoeuvre moved out of the way to pass the slower pedestrian in front, who stayed close to its preferred path. When the risk-thresholds were different, the high risk-threshold pedestrian still deviated the most, where in six cases the low risk-threshold pedestrian also made some effort to create space. The framework with the naïve belief lacks this behaviour, where with the same risk-thresholds both the pedestrian in front, as well as the one in the back deviated from their paths during the overtake. The pedestrian in front deviated even more than the one performing the overtaking manoeuvre, who in four cases barely deviated at all. When the risk-thresholds were different it was shown that the orange pedestrian yielded earlier and more in all cases, and the blue overtaking pedestrian only made any real effort to deviate from its preferred path in three cases.

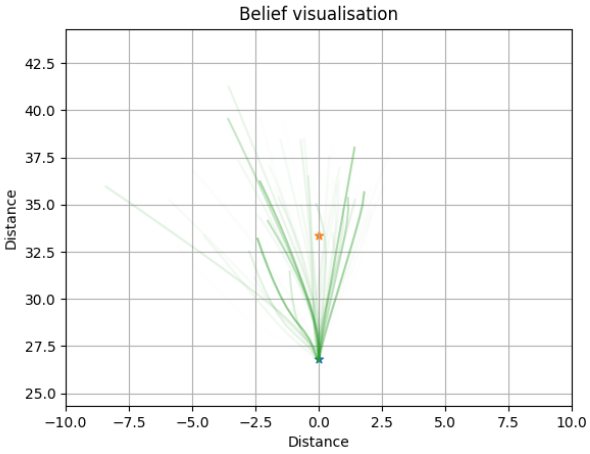


(a) TrajFlow belief

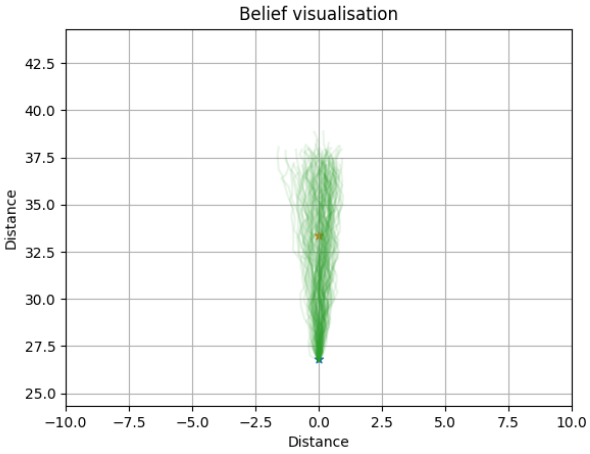


(b) Naïve belief

Figure 15: Visualisation of TrajFlow belief 15a and naïve belief 15b in the overtake scenario.

### D. Same goal scenario

*1) Expectations:* It is expected that both pedestrians will walk toward each other before a conflict arises after which they will yield to avoid each other until one of the pedestrians reaches the destination. Since this scenario is symmetric, the expectation for pedestrians with equal risk-thresholds is that both pedestrians yield an equal amount for each other on average. Note that the belief construction is stochastic, so variations are expected. When the pedestrians have different risk-thresholds it is expected that the low risk-threshold pedestrian yields earlier and more, allowing the high risk-threshold pedestrian to yield less and reach the goal. In both cases, the magnitude and direction of yielding will depend on the shape of the constructed belief.

*2) Observations:* Visualisations of the examples are shown in figure 16. Let's begin by examining the figures 16a and 16e where the risk-thresholds were the same for both pedestrians and the belief was created with TrajFlow. Conforming to the expectation, the pedestrians walked towards each other, before yielding a comparable amount, after which one of the pedestrians proceeded towards the goal, keeping the other at a distance. Looking at the rest of the simulations (Appendix B.4), similar behaviour emerged. In five cases one of the pedestrians reached the goal quickly, while in the other five cases, they first spiraled around the goal before one of the pedestrians reached it.

Next, we'll examine the figures 16b and 16f, where the

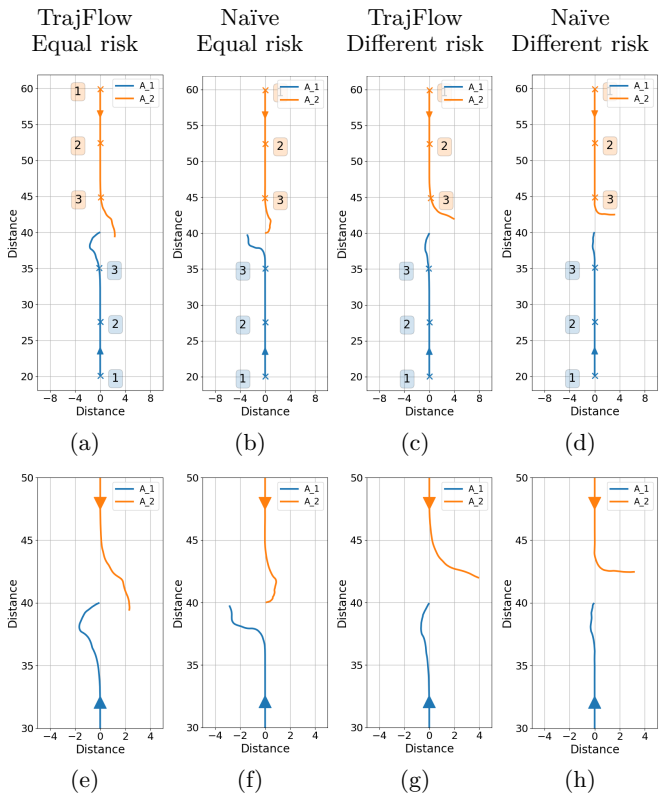


Figure 16: Same goal scenario trajectory examples. 16a uses TrajFlow belief construction and has equal risk-thresholds, 16b uses the naïve belief and has equal risk-thresholds, 16c uses TrajFlow belief and has different risk-thresholds, 16d uses the naïve belief and has different risk-thresholds. When the risk-thresholds are different, the orange pedestrian has lower risk-thresholds. Underneath in 16e, 16f, 16g, 16h the same scenarios are shown zoomed in on the interaction with equal x-axis and y-axis proportions.

risk-thresholds were also the same for both pedestrians, but the naïve belief construction was used. The trajectories look similar to the ones created with TrajFlow, where both pedestrians walked toward each other before yielding, followed by one pedestrian reaching the goal. A noticeable difference was that with the naïve belief the pedestrians yielded later and more drastic. The other simulations (Appendix B.4) reflect this behaviour, showing both pedestrians walking towards the goal before both yielding, followed by one pedestrian reaching the goal. In seven of the cases the pedestrian not reaching the goal was moving backwards at some point.

For the third configuration, figure 16c and 16g, the risk-thresholds were different for both pedestrians and the belief was created with TrajFlow. The behaviour was as expected, showing that the low risk-threshold pedestrian yielded earlier and more drastically, allowing the high risk-threshold pedestrian to yield less and reach the goal. Looking at the other nine trajectories, this behaviour was reflected in all cases (Appendix B.4). The metrics also show that the high risk-threshold pedestrian had a significantly lower maximum deviation, and the low risk-threshold pedestrian yielded significantly more and earlier, all compared to situations with equal risk-thresholds, as shown in figure 17a, 17b and 17c.

For the last configuration, figure 16d and 16h, the risk-thresholds were different for both pedestrians and the naïve belief construction was used. The trajectories had a similar shape to the ones with TrajFlow belief, where the high risk pedestrian stayed closer to its preferred path and reached the goal. Looking at the other simulations (Appendix B.4), similar behaviour was shown, with a difference compared to TrajFlow in the sharpness of the turns when yielding. In one case the goal is reached by the low risk-threshold pedestrian, but on average the low risk pedestrian deviated more, which is also shown as a significant increase in the maximum deviation metric shown in figure 17d.

*3) Beliefs:* The observed behaviour described above can be attributed to the difference in belief construction. Examples of the data-driven TrajFlow belief and the naïve belief are shown in figure 18. The belief created by TrajFlow, where beliefs with a higher likelihood are plotted in darker green, shows a wider spread in beliefs compared to the more narrow naïve beliefs. A consequence of this is that in the simulations containing the naïve belief, the perceived risk will stay low for longer before suddenly increasing quite sharply, explaining the late and more abrupt deviations of trajectories. A reason for the occurring spiralling behaviour with the data-driven belief is that there is more variation between beliefs at different time points, this can cause an oscillation where both pedestrians think it is safe to move towards the goal, before creating a new belief pushing their perceived risk above the threshold. In the naïve belief the spread is more consistent in the direction of movement, so once the stochasticity enables one pedestrian to take control the other pedestrian will not change its direction towards the goal, as this would push their perceived risk over the threshold.

*4) Overall influence TrajFlow:* Visually comparing the trajectories of the framework with TrajFlow and naïve beliefs in this scenario demonstrates a slight preference for using TrajFlow over the naïve belief implementation, although both do not show perfect behaviour. When the risk-thresholds were equal, in half of the cases the trajectories showed plausible interaction behaviour where both pedestrians yielded to avoid each other, after which one pedestrian found the way to the goal, and the other kept moving around it in a curve. In the other half of the cases, both pedestrians got stuck in a spiral

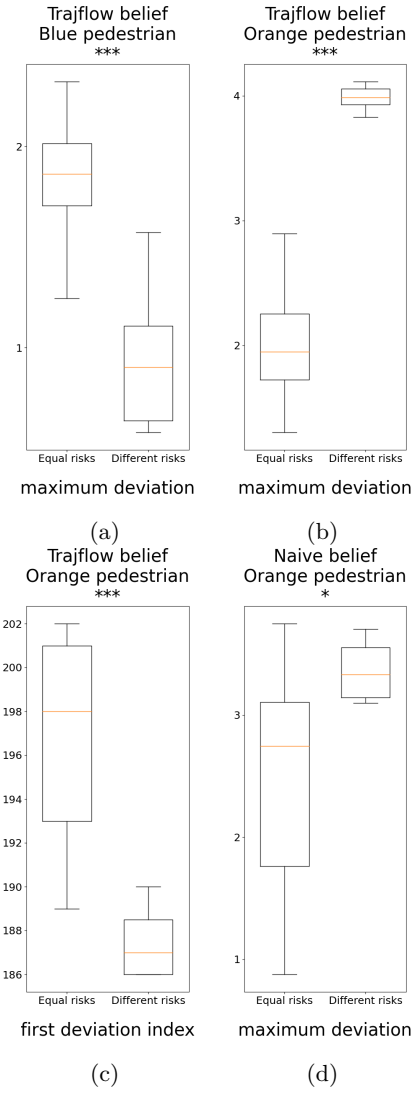


Figure 17: Boxplots comparing metrics of pedestrians between equal and different risk-thresholds in the same goal scenario.

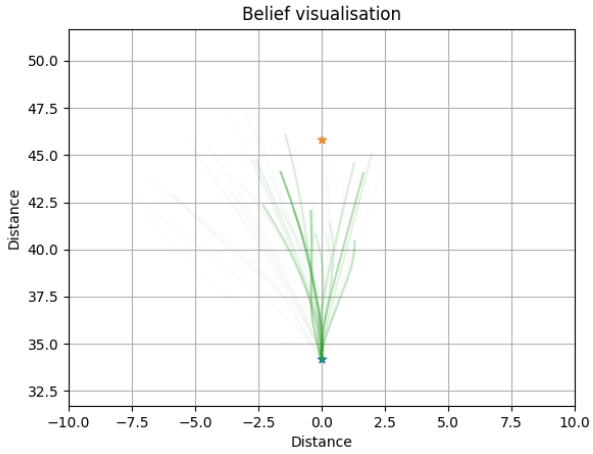


(a) TrajFlow belief

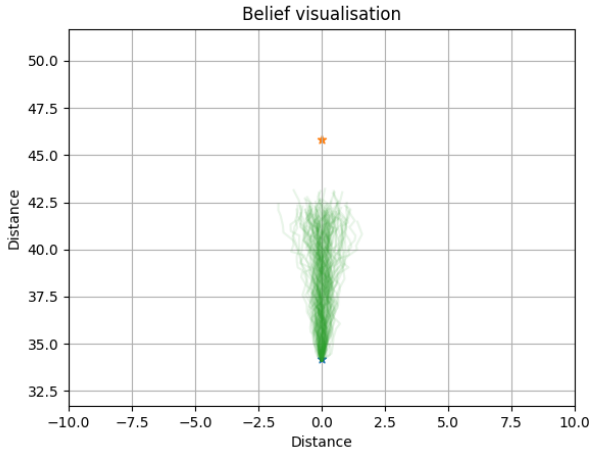


(b) Naïve belief

Figure 18: Visualisation of TrajFlow belief 18a and naïve belief 18b in the same goal scenario.

around the goal before one of them reached it. When the risk-thresholds were different the low risk-threshold pedestrian deviated earlier and more, enabling the high threshold pedestrian to deviate less and reach the goal in all cases. The framework with naïve beliefs and equal risk-thresholds performs better because it does not show the spiraling behaviour that TrajFlow had, but also worse because in seven scenarios, one pedestrian is moving backwards while the other one reaches the goal. When the risk-thresholds were different the low risk-threshold pedestrian showed a very sharp change in trajectory, and the high threshold agent could use that most of the time to get to the goal. In one case the low risk-threshold pedestrian reached the goal while the high risk-threshold pedestrian was walking away.

### *E. Left-side bias*

In the simulations with TrajFlow belief where pedestrians were walking towards each other, like the straight and same goal scenario, a pattern emerges: the pedestrians consistently moved to their left in order to avoid each other. To figure out why that is the case, the beliefs of one pedestrian walking towards another were plotted over multiple timesteps as can be seen in figure 19. This plot indicates that the TrajFlow belief has a higher probability of predicting trajectories that are going to the left. A consequence is that the other agent will create plans that pass on the other (right) side to keep the perceived risk lower. There are several reasons which might cause this that will be discussed in section V.

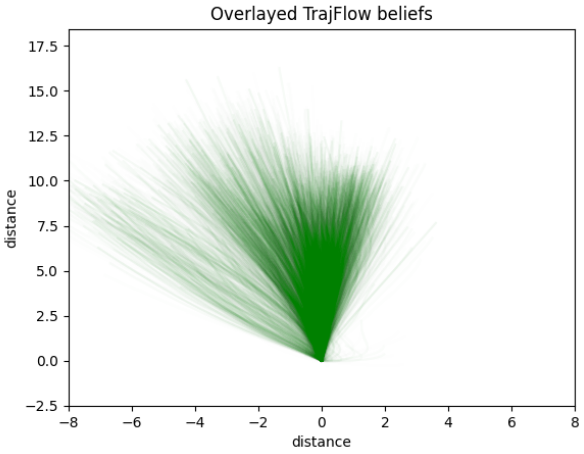


Figure 19: Visualisation of overlayed TrajFlow beliefs from multiple timesteps.

## V. Discussion

The objective of this research is to find an answer to the research question:

What influence does the integration of TrajFlow's multimodal trajectory prediction in the Communication-Enabled-Interaction framework have on simulated pedestrian behaviour in interactive traffic scenarios?

To investigate this, the Communication-Enabled-Interaction framework has been extended to work for pedestrians in two-dimensional scenarios. Modifications in the overall dynamics, planning, and risk evaluation were made to accommodate this. TrajFlow was trained on the Forking Paths dataset and has been integrated into the framework so that its trajectory predictions could be utilised for the belief construction in the framework. Four interactive pedestrian simulation scenarios were created in which the influence of TrajFlow's integration could be investigated both qualitatively and quantitatively.

In most cases, using the data-driven TrajFlow belief construction has an advantageous effect on the realism of the simulated pedestrian behaviour. The plausible interaction behaviour is especially noticeable in the straight and overtake scenario, where the beliefs show a clear indication of the expectation to go left or right instead of continuing straight. In the straight scenario, the pedestrians take smoother paths to pass each other, and in the overtake scenario the belief that the overtaking pedestrian would deviate from its path enabled the pedestrian that was being overtaken to stay on its preferred path. In the cross and same-goal scenario, the advantage of using TrajFlow is less obvious. In the cross scenario, the behaviour shows mostly plausible interactions and the pedestrians do not deviate far from their path to let the other one cross, however, there are exceptions where the conflict resolution results in odd-looking trajectories. Compared to the naïve belief construction it does show more promise, as there the conflict resolution often spanned a long time, causing the pedestrians to take sub-optimal paths towards their goals. In the same-goal scenario, when the risk-thresholds are different, the expected behaviour is shown in the simulations using TrajFlow. When the risk-thresholds are equal, however, in half the cases the pedestrians got stuck in a spiral before one of them reached the goal.

This brings us to the limitations of using the data-driven TrajFlow belief in the CEI framework and the limitations of this study as a whole.

### A. Limitations

*1) Simulations:* A limitation of this work was the limited number of simulations that have been run. While the original CEI framework is deterministic and would have required only one simulation run for each scenario, the variability introduced by the belief construction made it necessary to conduct multiple simulations for each specific scenario and setup. However, these simulations were time-consuming, so a trade-off had to be made between running as many simulations as possible and spending a reasonable amount of time. In the end, the decision was made to run 10 simulations per scenario, which together with 16 different scenario setups amounted to 160 simulations. However, this finite number is not enough to fully cover the entire spectrum of possible events. As a consequence, the data might not have accurately reflected the complete range of pedestrian behaviours and interactions.

A consequence of the limited runs of simulations was the limited variation in risk-thresholds and initial simulation conditions. The simulations were performed using a set of distinct risk-thresholds (equal or different) and a single set of initial conditions per scenario. While being sufficient for a proof of concept, these did not cover the entire spectrum of possibilities. The true risk-thresholds of pedestrians cover a wide range of risk taking behaviours influenced by factors such as personality, past experiences and situational context. Using a wider range of risk-thresholds would have provided a more nuanced understanding of how these risk-thresholds impact the interaction modelling. The initial simulation conditions were a predefined set of pedestrian placements, velocities and goal positions. This limited the span of possible interactions and emerging behaviours that were explored in each scenario. The use of a variety of initial conditions could have accommodated a better analysis of the robustness and variability of the pedestrian behaviour.

Another aspect to consider is how accurately our simulations reflect real-life situations. In the overtake scenario for example, the framework provided the pedestrian walking in front with indirect communication about the velocity and position of the pedestrian behind.

When getting overtaken you usually do not have this information. In the same goal scenario, in the real world when you see someone going exactly where you are going, you use visual and social clues to infer whether that person is going to stay there shortly so you can wait and temporarily change your goal, or find a new final goal all together. In the case that you really would want to go to that exact position where someone else is going, you would start a conversation about it, all aspects that were not possible in the simulations.

These simplified scenario representations formed a limitation. Although the aim of the created pedestrian scenarios was to mirror real-world interactions like collision avoidance, overtaking, and shared goals, their design in a simulation environment is inherently simplified. Actual pedestrian dynamics in interactions involve a great many factors that were beyond the scope of the scenario simulations such as terrain and obstacles, diverse behaviours like pedestrian group dynamics, and context factors such as crowded spaces which all play a role in shaping pedestrian interactions. The simulation scenarios may therefore have oversimplified the real-world complexities and potentially limited the generalizability of this study's findings.

*2) Training TrajFlow:* A limitation that might be in the training of TrajFlow is the representativeness of the training data. The predictive capabilities of TrajFlow were influenced by the data that it was trained on because it relied on learning the patterns and behaviours that are present in the training dataset. The dataset might have introduced biases inherent to that set. If, for instance, the dataset did not capture certain pedestrian behaviours, TrajFlow might not have been able to generalize well for these types of behaviours.

Another limitation was the way the data was processed before the learning process started. For all the scenarios from the data, only the eight timesteps leading up to the branching in multi-future scenarios, and the twelve timesteps after, were considered. This might restrict the contextual understanding capabilities and limit the anticipation of long-term trajectories or interactions formed over extended periods of time.

This brings us to a limitation that came from the Forking Paths dataset. While this dataset is grounded in real-world data, it provides multiple plausible future trajectories that are created by a limited number of ten annotators. This group might not have been a good representation of the entirety of pedestrian interaction behaviour and might not contain outlier behaviour, or indeed too much of it, limiting the representativeness of TrajFlow's belief construction and therefore the capacity of the modelled pedestrian interactions.

Moreover, the dataset contains 750 sequences across 7 scenes, but after removing the sequences containing interactions with vehicles, this was reduced to 184 sequences across 4 scenes. The dataset was therefore limited in scope, possibly lacking variations in interactions present in crowded or unconventional settings, and possibly restrained the belief construction to a specific subset of scenarios. If most of the trajectories in the dataset were pedestrians moving to their left when they avoid someone, that behaviour would be reflected in the trajectories that TrajFlow would predict once trained on this data. This might very well be the reason why the pedestrians using the data-driven TrajFlow belief often evade each other on the left.

Since the different scenes in the dataset were limited, it is also very likely that there are scenarios that do not occur often. If for example there are almost no instances of pedestrians crossing paths diagonally, it can not be expected that TrajFlow can accurately predict how pedestrians act in those situations.

*3) Framework integration:* Another potential limitation was in the mismatch of the CEI timestep $\Delta t = 0.05s$ and TrajFlows timestep $\Delta t = 0.4s$. This might cause difficulties in accurately capturing nuanced and quick changes in behaviour or trajectory variations. Consequently, this could limit the effectiveness of TrajFlows prediction in the "real-time" CEI decision-making context. Furthermore, the filtering of the position plan timesteps to only match the ones of the belief for risk evaluation leads to a significant loss of information. This can especially become problematic in scenarios where the velocity over the last timesteps was low and increases suddenly. One might think an easy fix is to either increase the timesteps in the CEI framework or decrease the timesteps in which TrajFlow operates. Unfortunately, this will introduce new problems. Increasing the timestep in which the CEI framework, and consequently the simulation, runs to $\Delta t = 0.4s$ will reduce the computational power needed to perform simulations, but could significantly compromise the accuracy and quality of those simulations. Finer details and rapid changes within the pedestrian behaviours might be overlooked or misrepresented due to the greater timestep. On the other hand, decreasing the timestep for TrajFlow predictions poses another problem. Trajflow was trained to use the last 8 timesteps to predict the following 12 steps in the future. With a timestep of $\Delta t = 0.4s$ and an average human walking speed, this translated to a distance of $4.8m$ and $7.2m$ respectively. Decreasing the timestep to $\Delta t = 0.05s$ would decrease the trajectory distance from which TrajFlow must learn to $0.6m$ and $0.9m$. The length of the trajectories would be too short to extract any meaningful information. To maintain the same predictive distance of $7.2m$ with a timestep $\Delta t = 0.05s$, TrajFlow has to produce trajectories of 96 points into the future. This is four times more than the maximum that TrajFlow has been tested on, and it might run into issues with compounding errors due to its autoregressive nature.

### B. Recomendations and further research

With regard to the gained insights and encountered limitations in this study, multiple recommendations for further research come to mind. These offer opportunities to expand and improve on this work and could offer valuable insights and advancements in modelling pedestrian behaviours within interactive traffic scenarios.

*1) Enhanced Dataset Diversity:* Further research should focus on the generation or acquisition of datasets that contain a broader spectrum of pedestrian behaviour in interactive scenarios, preferably as an extension to the Forking Paths dataset. These datasets could then be used in training to improve the predictive capabilities of TrajFlow with more behavioural patterns to see whether substantial changes in behaviour emerge.

*2) Real-World Complexity Integration:* Another valuable course of action to continue this research would be to investigate ways that integrate more real-world complexity into pedestrian simulation scenarios. This could include obstacles, group dynamics, crowded spaces and varying terrains (pavement/road/grass), which will significantly impact the interactions.

*3) Dynamic Risk Perception and Responses:* It would be very interesting to study mechanisms that enable dynamic risk-thresholds. Investigating how humans change their risk-thresholds based on the current circumstances and actions of others could significantly influence interaction modelling.

*4) Temporal Alignment Enhancement:* Research addressing the mismatch between TrajFlow's predictions timestep and CEI's simulation timestep could help optimize the integration of trajectory predictions within real-time decision-making frameworks.

*5) Multi-Agent Interactions and Communication:* Using the CEI framework for multi-agent communication and interaction would enable the investigation of interesting pedestrian interaction scenarios where its workings can be examined in scenarios guided by group dynamics.

*6) Benchmarking and Generalizability Testing:* Conducting an extensive benchmarking investigation on various integrated trajectory prediction models with different interactive pedestrian scenarios and datasets. This analysis would offer insight into the robustness of models and performance across conditions.

*7) Human-Robot Interaction Studies:* Exploring pedestrian behaviour modelling in the context of interactions with autonomous agents or robots. Studying how pedestrians react and adjust their behaviour when facing interactions with automated systems and comparing the similarities and differences with pedestrian-pedestrian behaviour could be a crucial development for safer traffic scenarios in the future.

## VI. Conclusion

In conclusion, this research tries to overcome the gap between sophisticated interaction modelling frameworks like the Communication-Enabled-Interaction framework and complex trajectory prediction models like TrajFlow by integrating them into a single framework focussed on the modelling of pedestrian interaction behaviour and testing it on four interactive scenarios.

The findings show that beliefs created by TrajFlow can give a significant boost in simulating plausible pedestrian interaction behaviour, but there are also some drawbacks concerning for example training data and timing-related issues. These limitations in simulation constraints and dataset biases emphasize the need for further research.

Nonetheless, a foundation is created on which future research can build and possible research directions are presented, aiding the enhancement of pedestrian behaviour modelling in interactive traffic scenarios by combining data-driven prediction models with interaction modelling frameworks, and contributing to safer urban environments and intelligent mobility systems.

## References


[1] J. M. Anderson, N. Kalra, K. D. Stanley, P. Sorensen, C. Samaras, and T. A. Oluwatola, *Autonomous Vehicle Technology: A Guide for Policymakers*. Santa Monica, CA: RAND Corporation, 2016.

[2] T. T. M. Tran, C. Parker, and M. Tomitsch, "A review of virtual reality studies on autonomous vehicle–pedestrian interaction," *IEEE Transactions on Human-Machine Systems*, vol. 51, no. 6, pp. 641–652, 2021.

[3] A. Der Kiureghian and O. Ditlevsen, "Aleatory or epistemic? does it matter?" *Structural Safety*, vol. 31, pp. 105–112, 03 2009.

[4] J. A. Michon, "A critical view of driver behavior models: What do we know," 1985. [Online]. Available: https://api.semanticscholar.org/CorpusID:15590886

[5] B. Varadarajan, A. Hefny, A. Srivastava, K. S. Refaat, N. Nayakanti, A. Cornman, K. Chen, B. Douillard, C. P. Lam, D. Anguelov, and B. Sapp, "Multipath++: Efficient information fusion and trajectory aggregation for behavior prediction," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 7814–7821.

[6] K. Messaoud, N. Deo, M. M. Trivedi, and F. Nashashibi, "Trajectory prediction for autonomous driving based on multi-head attention with joint agent-map representation," in *2021 IEEE Intelligent Vehicles Symposium (IV)*, 2021, pp. 165–170.

[7] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social gan: Socially acceptable trajectories with generative adversarial networks," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2255–2264.

[8] J. Amirian, J.-B. Hayet, and J. Pettré, "Social ways: Learning multi-modal distributions of pedestrian trajectories with gans," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2019, pp. 2964–2972.

[9] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *Computer Vision – ECCV 2020*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds. Cham: Springer International Publishing, 2020, pp. 683–700.

[10] Y. Yuan, X. Weng, Y. Ou, and K. Kitani, "Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting," in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 9793–9803.

[11] B. Brito, H. Zhu, W. Pan, and J. Alonso-Mora, "Social-vrnn: One-shot multi-modal trajectory prediction for interacting pedestrians," 2020.

[12] A. Bertugli, S. Calderara, P. Coscia, L. Ballan, and R. Cucchiara, "Ac-vrnn: Attentive conditional-vrnn for multi-future trajectory prediction," *Computer Vision and Image Understanding*, vol. 210, p. 103245, Sep. 2021. [Online]. Available: http://dx.doi.org/10.1016/j.cviu.2021.103245

[13] A. Mészáros, J. F. Schumann, J. Alonso-Mora, A. Zgonnikov, and J. Kober, "Trajflow: Learning the distribution over trajectories," 2023.

[14] S. Pellegrini, A. Ess, K. Schindler, and L. van Gool, "You'll never walk alone: Modeling social behavior for multi-target tracking," in *2009 IEEE 12th International Conference on Computer Vision*, 2009, pp. 261–268.

[15] A. Lerner, Y. Chrysanthou, and D. Lischinski, "Crowds by Example," *Computer Graphics Forum*, 2007.

[16] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese, "Learning social etiquette: Human trajectory understanding in crowded scenes," in *European Conference on Computer Vision*, 2016. [Online]. Available: https://api.semanticscholar.org/CorpusID:3150075

[17] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscenes: A multimodal dataset for autonomous driving," 2020.

[18] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: the kitti dataset," *The International Journal of Robotics Research*, vol. 32, pp. 1231–1237, 09 2013.

[19] J. Liang, L. Jiang, K. Murphy, T. Yu, and A. Hauptmann, "The garden of forking paths: Towards multi-future trajectory prediction," in *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

[20] G. Markkula, R. Madigan, D. Nathanael, E. Portouli, Y. M. Lee, A. Dietrich, J. Billington, A. Schieben, and N. Merat, "Defining interactions: a conceptual framework for understanding interactive behaviour in human and automated road traffic," *Theoretical Issues in Ergonomics Science*, vol. 21, 02 2020.

[21] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Phys. Rev. E*, vol. 62, pp. 1805–1824, Aug 2000. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevE.62.1805

[22] A. Kesting, M. Treiber, and D. Helbing, "Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 368, no. 1928, pp. 4585–4605, 2010. [Online]. Available: https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2010.0084

[23] M. Rahman, M. Chowdhury, Y. Xie, and Y. He, "Review of microscopic lane-changing models and future research opportunities," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 4, pp. 1942–1956, 2013.

[24] D. D. Salvucci and A. Liu, "The time course of a lane change: Driver control and eye-movement behavior," *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 5, no. 2, pp. 123–132, 2002. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1369847802000116

[25] A. Zgonnikov, D. Abbink, and G. Markkula, "Should i stay or should i go? cognitive modeling of left-turn gap acceptance decisions in human drivers," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, p. 001872082211445, 12 2022.

[26] K. Tian, G. Markkula, C. Wei, Y. M. Lee, R. Madigan, N. Merat, and R. Romano, "Explaining unsafe pedestrian road crossing behaviours using a psychophysics-based gap acceptance model," *Safety Science*, vol. 154, p. 105837, 10 2022.

[27] O. Siebinga, A. Zgonnikov, and D. A. Abbink, "Modelling communication-enabled traffic interactions," 2023.

[28] R. Tian, S. Li, N. Li, I. Kolmanovsky, A. Girard, and Y. Yildiz, "Adaptive game-theoretic decision making for autonomous vehicle control at roundabouts," in *2018 IEEE Conference on Decision and Control (CDC)*, 2018, pp. 321–326.

[29] Q. Zhang, D. Filev, H. E. Tseng, S. Szwabowski, and R. Langari, "Addressing mandatory lane change problem with game theoretic model predictive control and fuzzy markov chain," in *2018 Annual American Control Conference (ACC)*, 2018, pp. 4764–4771.

[30] S. Coskun, Q. Zhang, and R. Langari, "Receding horizon markov game autonomous driving strategy," in *2019 American Control Conference (ACC)*, 2019,

pp. 1367–1374.
[31] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," 2017.

# II

# Appendices

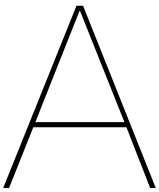

# Simulation parameters

| Straight path experiment | | | | |
|---|---|---|---|---|
| **Experiment** | **Belief** | **Risk agent 1** | **Risk agent 2** | **Runs** |
| 1.1 | Trajflow | 0.6 0.8 | 0.6 0.8 | 10 |
| 1.2 | Trajflow | 0.6 0.8 | 0.3 0.5 | 10 |
| 1.3 | Naïve | 0.7 0.9 | 0.7 0.9 | 10 |
| 1.4 | Naïve | 0.7 0.9 | 0.4 0.6 | 10 |

| Crossing path experiment | | | | |
|---|---|---|---|---|
| **Experiment** | **Belief** | **Risk agent 1** | **Risk agent 2** | **Runs** |
| 2.1 | Trajflow | 0.6 0.8 | 0.6 0.8 | 10 |
| 2.2 | Trajflow | 0.6 0.8 | 0.3 0.5 | 10 |
| 2.3 | Naïve | 0.7 0.9 | 0.7 0.9 | 10 |
| 2.4 | Naïve | 0.7 0.9 | 0.4 0.6 | 10 |

| Overtake experiment | | | | |
|---|---|---|---|---|
| **Experiment** | **Belief** | **Risk agent 1** | **Risk agent 2** | **Runs** |
| 3.1 | Trajflow | 0.6 0.8 | 0.6 0.8 | 10 |
| 3.2 | Trajflow | 0.6 0.8 | 0.3 0.5 | 10 |
| 3.3 | Naïve | 0.7 0.9 | 0.7 0.9 | 10 |
| 3.4 | Naïve | 0.7 0.9 | 0.4 0.6 | 10 |

| Same goal experiment | | | | |
|---|---|---|---|---|
| **Experiment** | **Belief** | **Risk agent 1** | **Risk agent 2** | **Runs** |
| 4.1 | Trajflow | 0.6 0.8 | 0.6 0.8 | 10 |
| 4.2 | Trajflow | 0.6 0.8 | 0.3 0.5 | 10 |
| 4.3 | Naïve | 0.7 0.9 | 0.7 0.9 | 10 |
| 4.4 | Naïve | 0.7 0.9 | 0.4 0.6 | 10 |

# B

# Qualitative Results

## B.1. Experiment: Straight

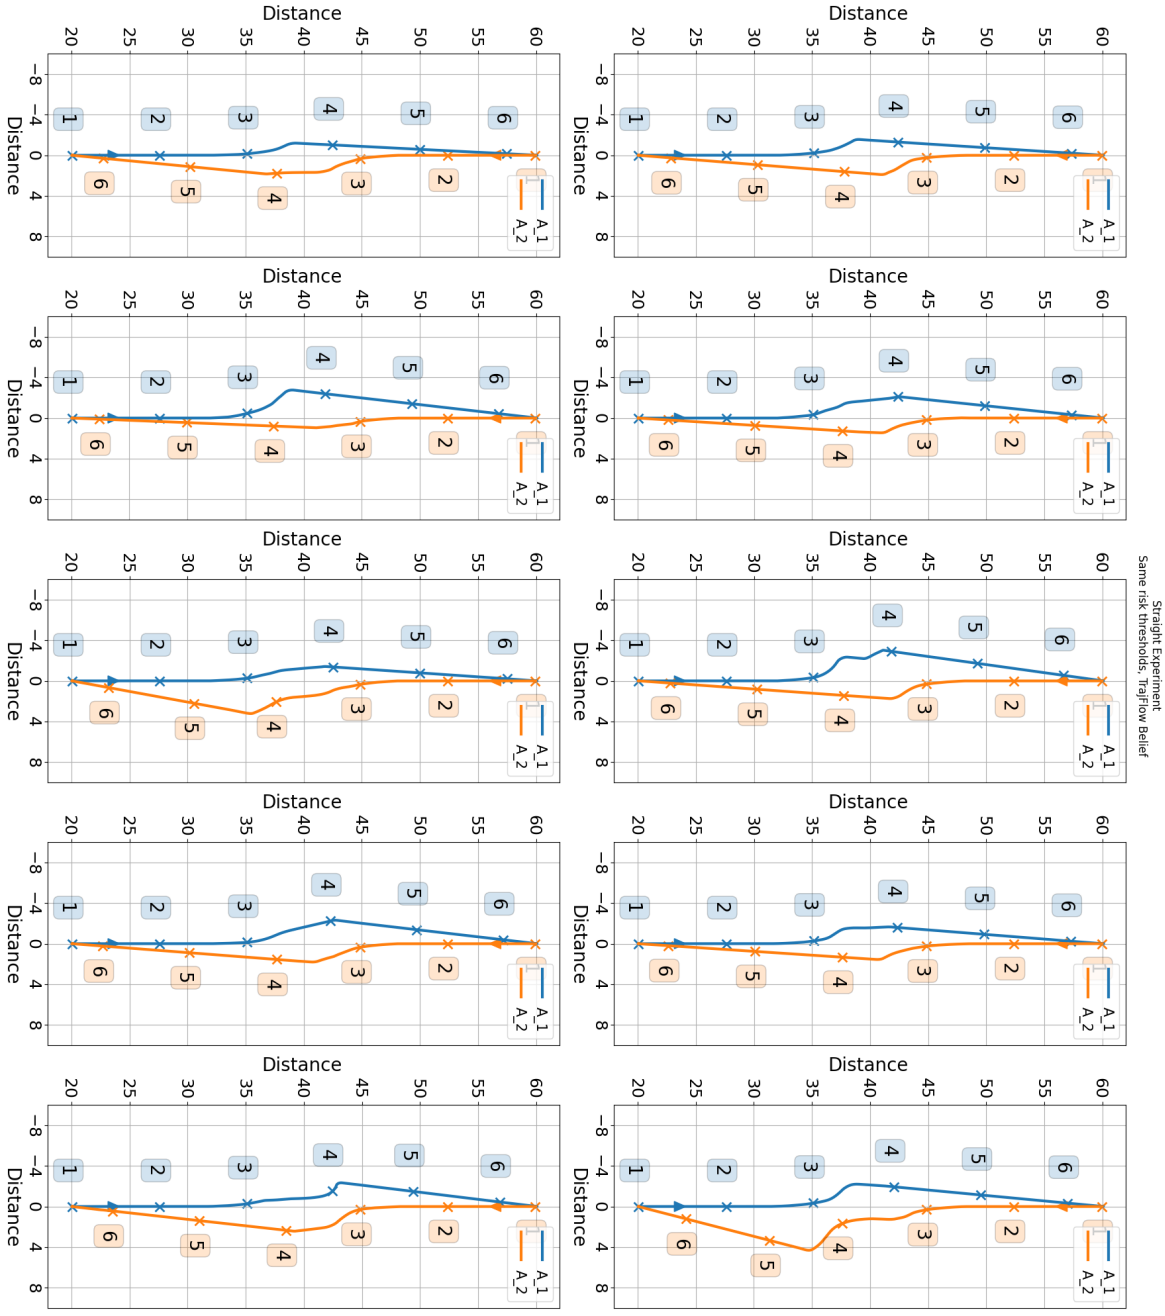


Figure B.1: Experiment: Straight, Risk threshold: Equal, Belief: TrajFlow

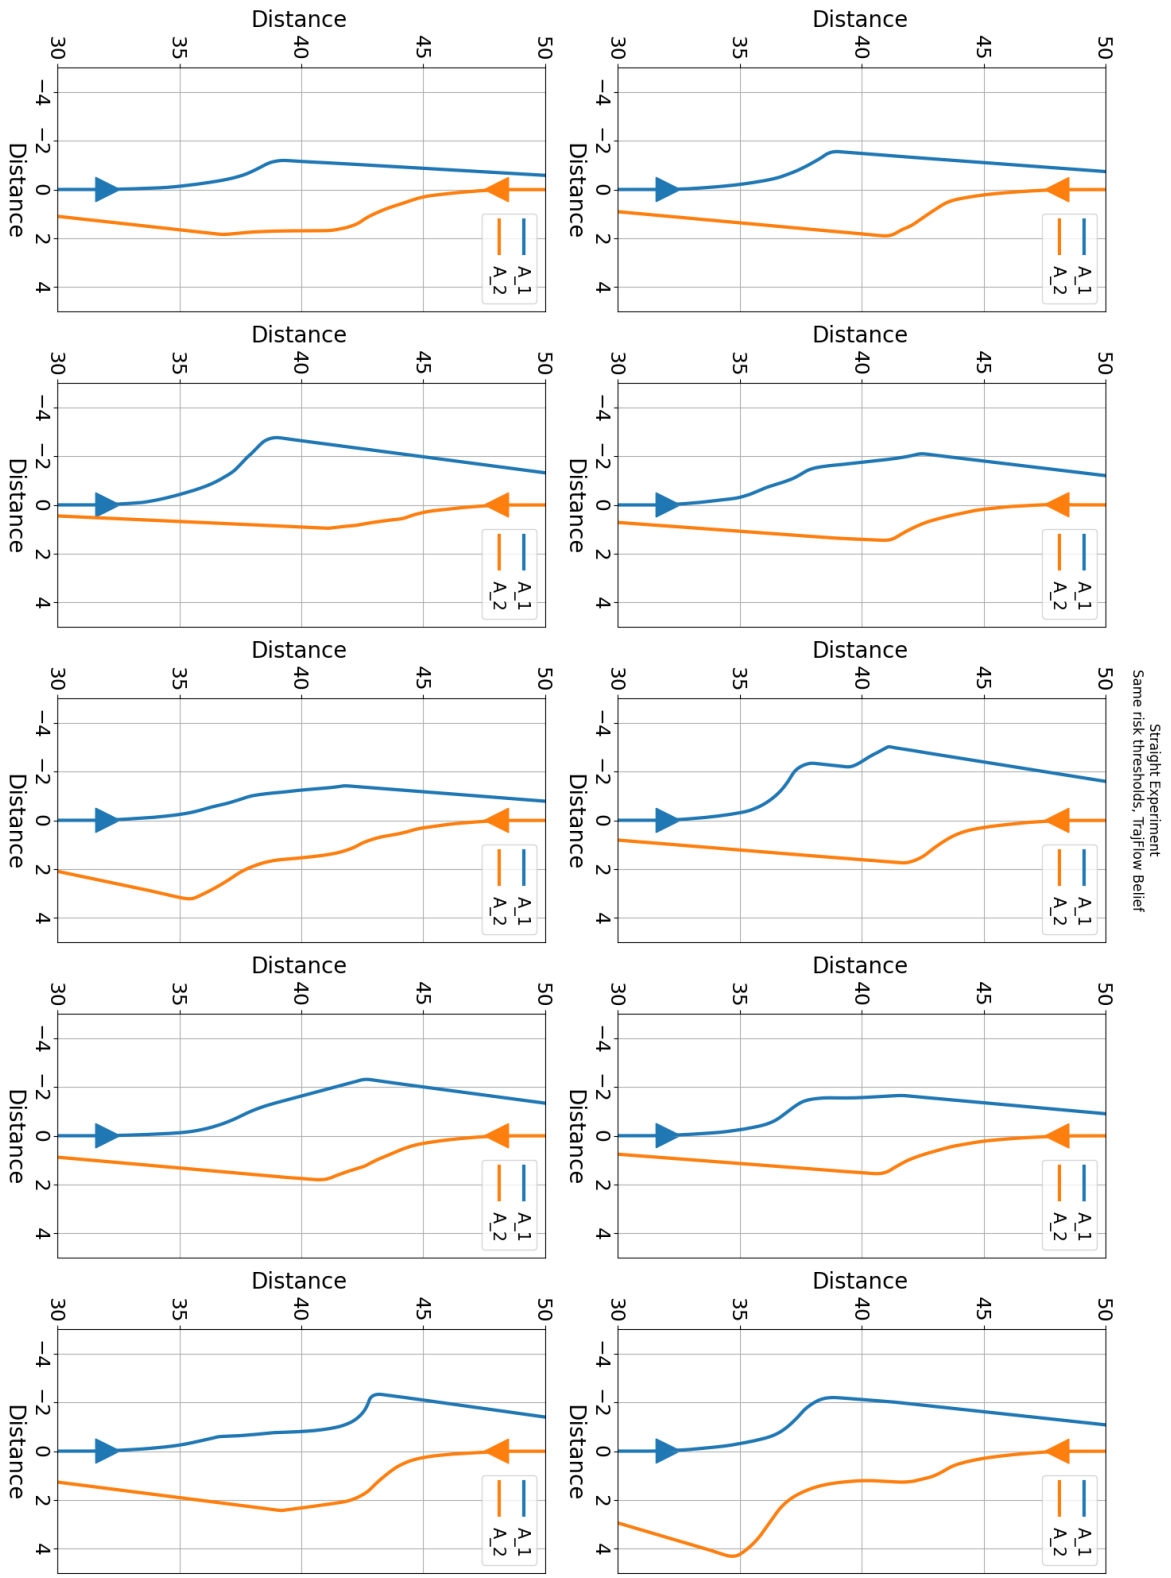


Figure B.2: Experiment: Straight, Risk threshold: Equal, Belief: TrajFlow, magnified

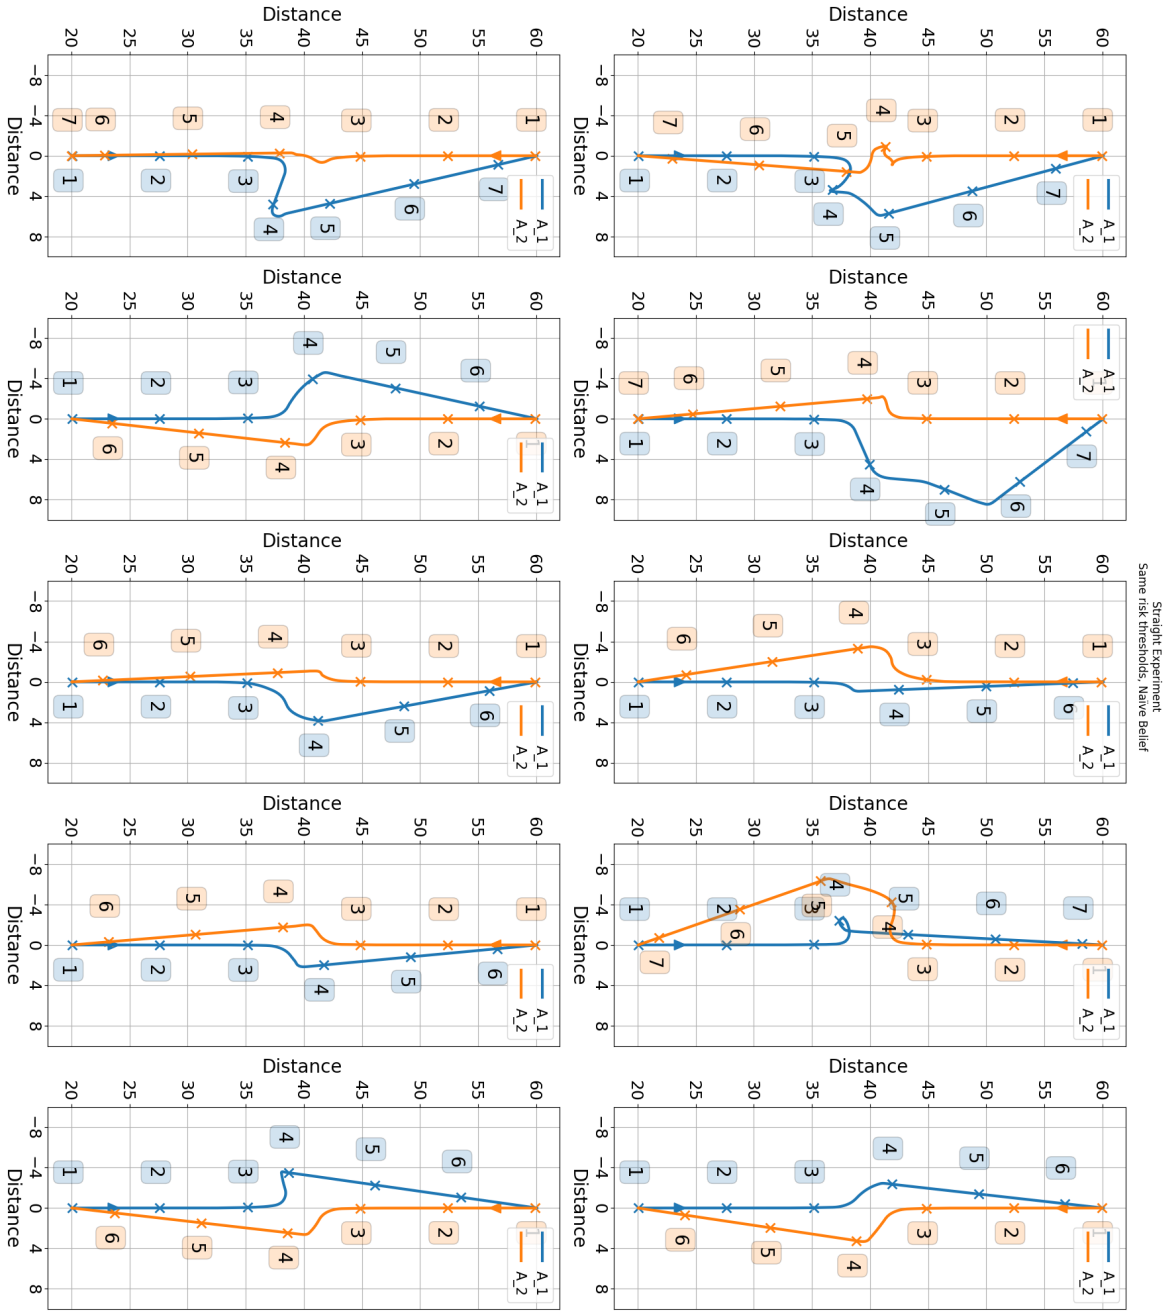


Figure B.3: Experiment: Straight, Risk threshold: Equal, Belief: Naïve

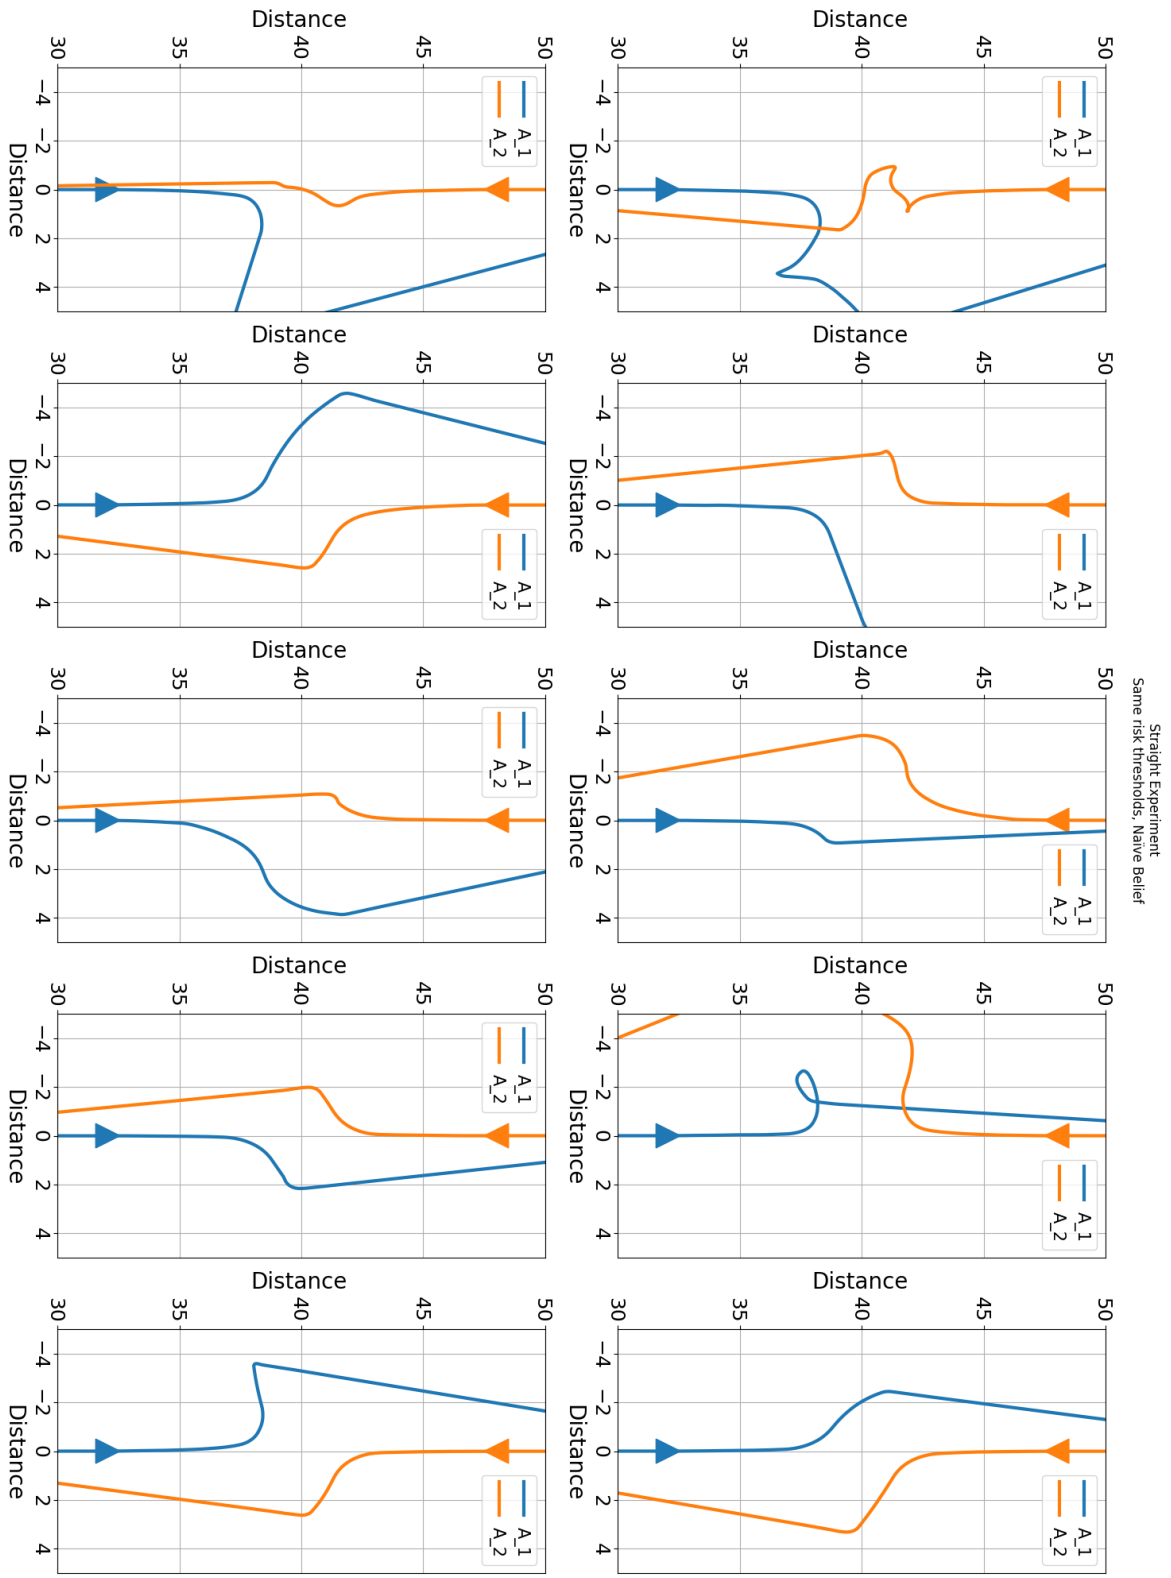


Figure B.4: Experiment: Straight, Risk threshold: Equal, Belief: Naïve, magnified

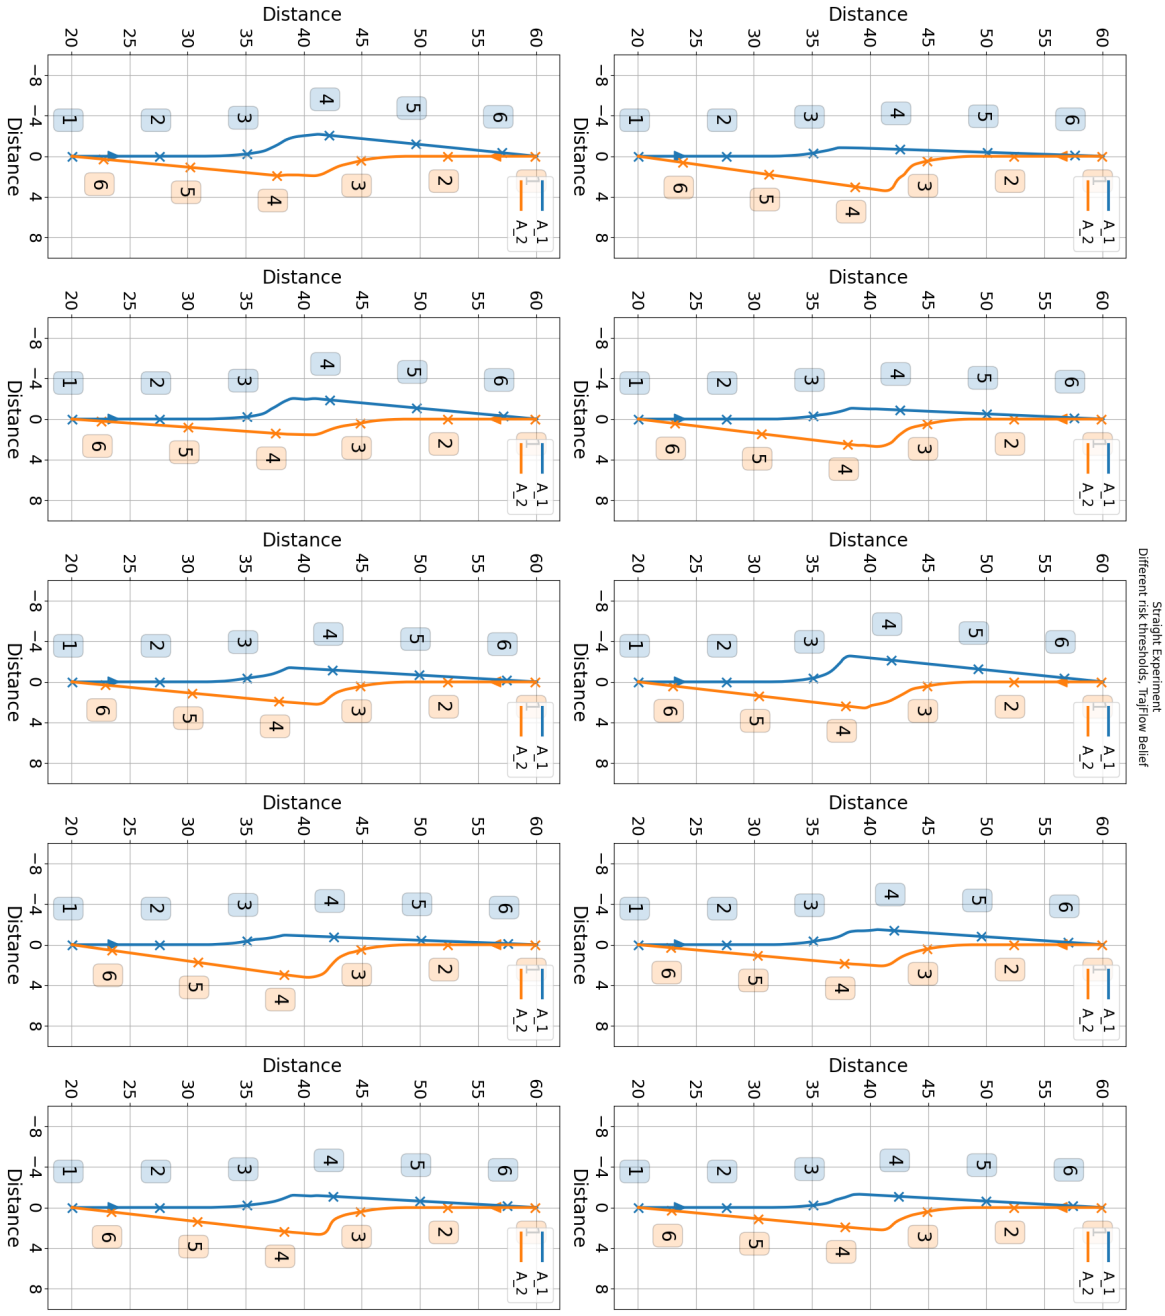


Figure B.5: Experiment: Straight, Risk threshold: Different, Belief: TrajFlow

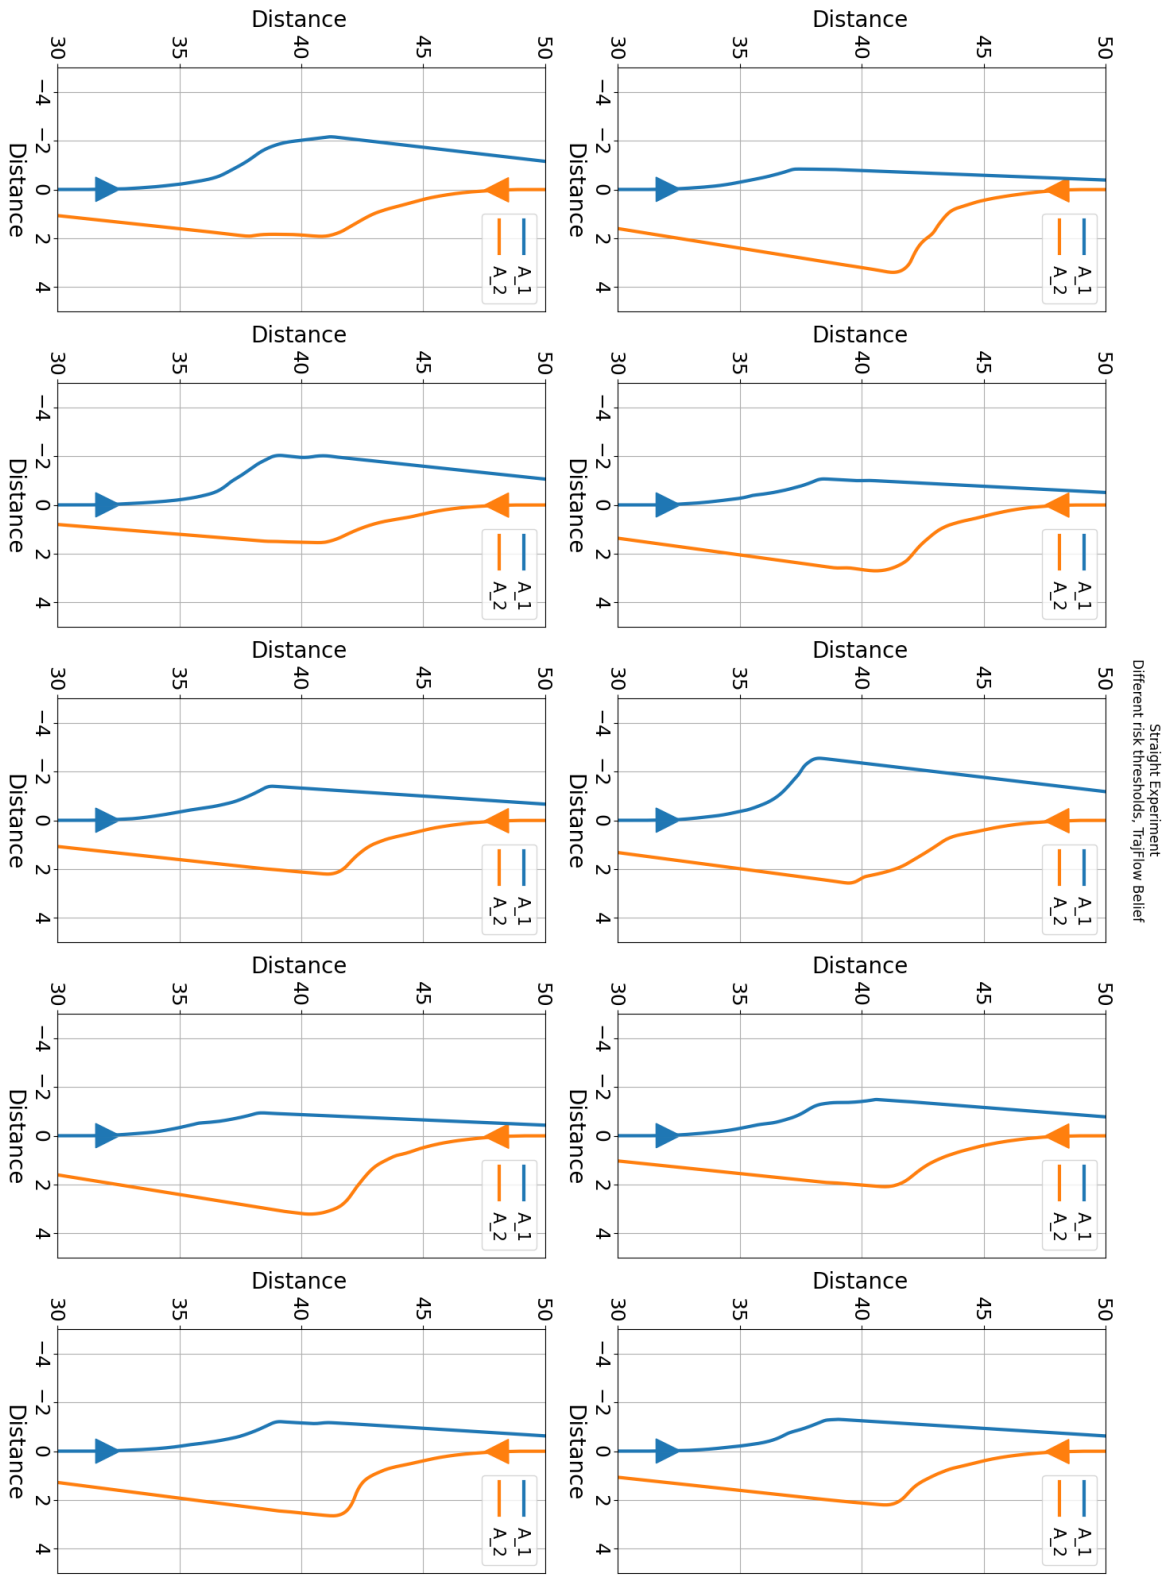


Figure B.6: Experiment: Straight, Risk threshold: Different, Belief: TrajFlow, magnified

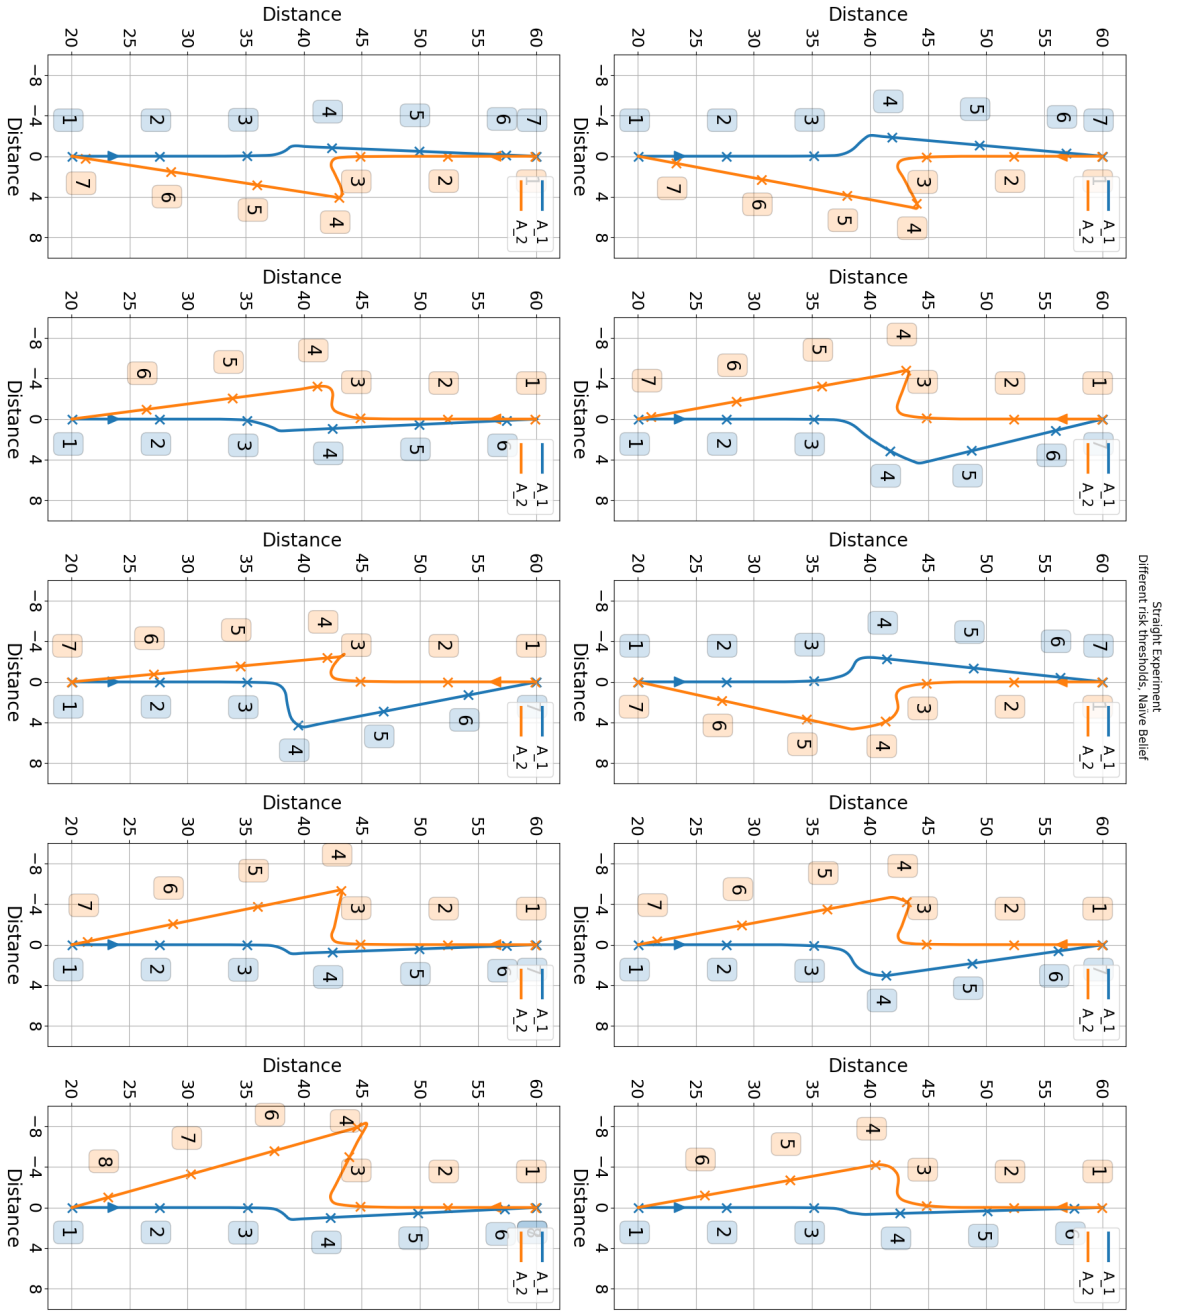


Figure B.7: Experiment: Straight, Risk threshold: Different, Belief: Naïve

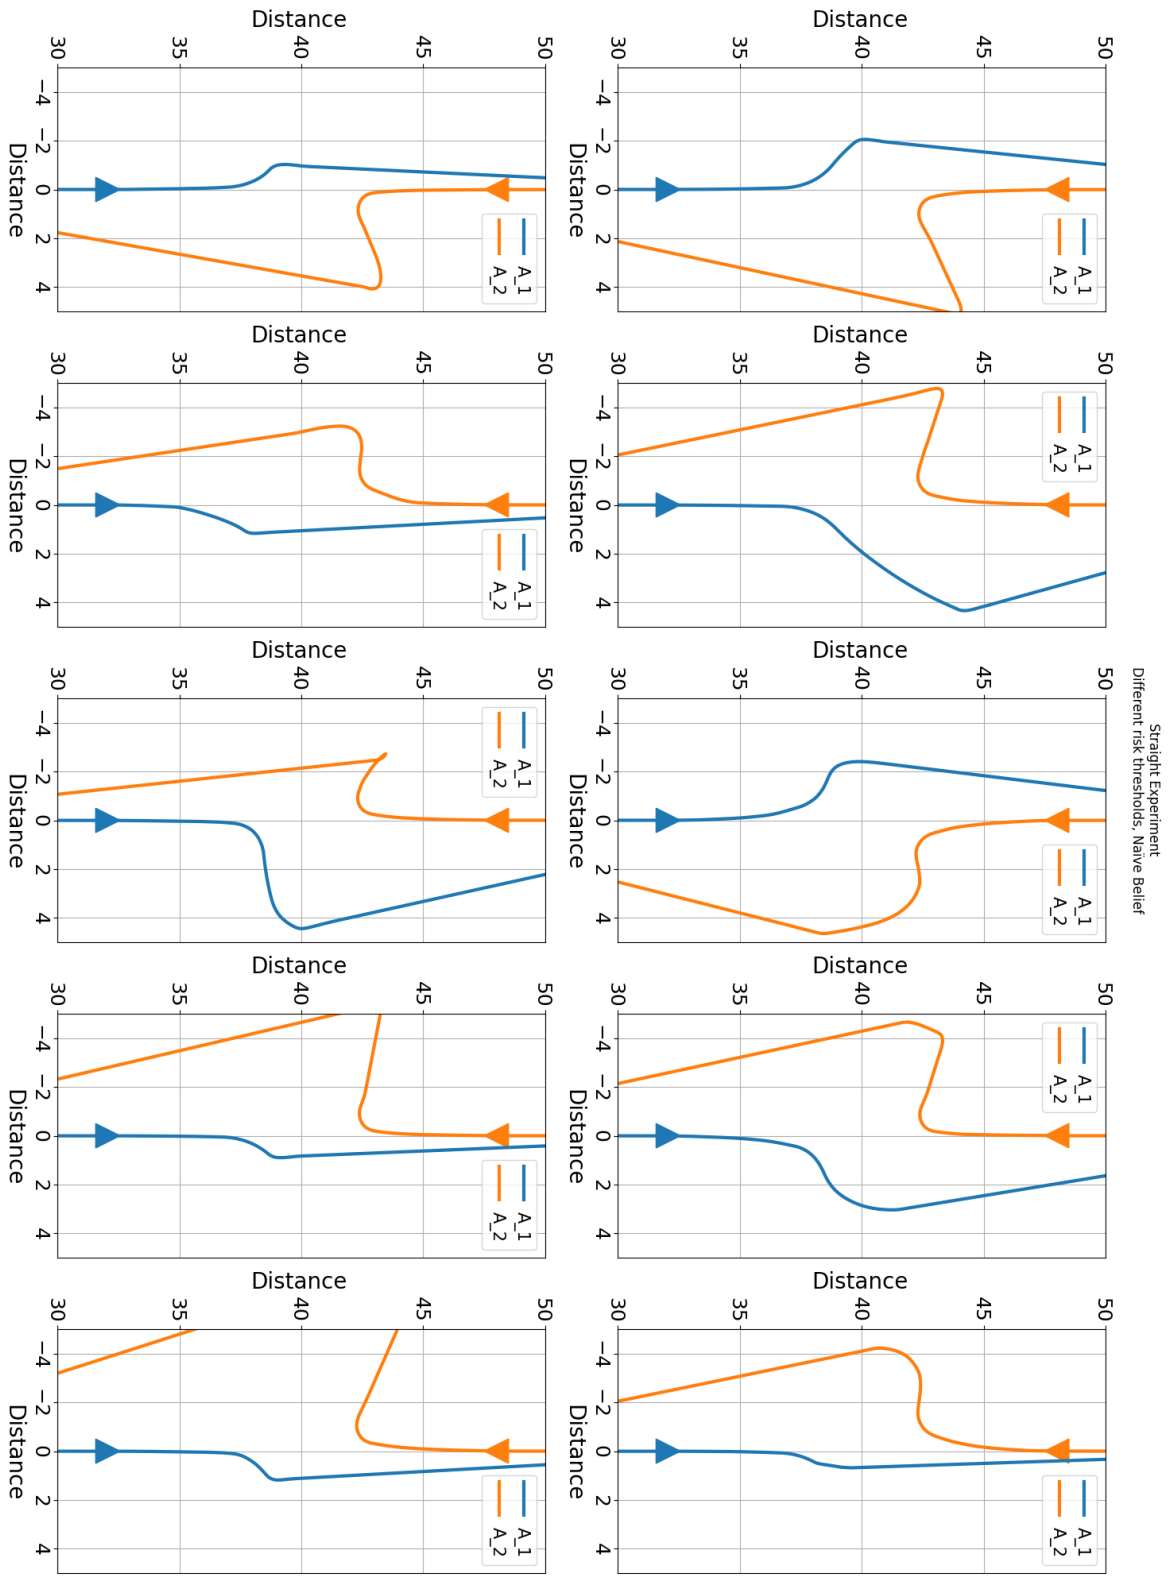


Figure B.8: Experiment: Straight, Risk threshold: Different, Belief: Naïve, magnified

## B.2. Experiment: Cross

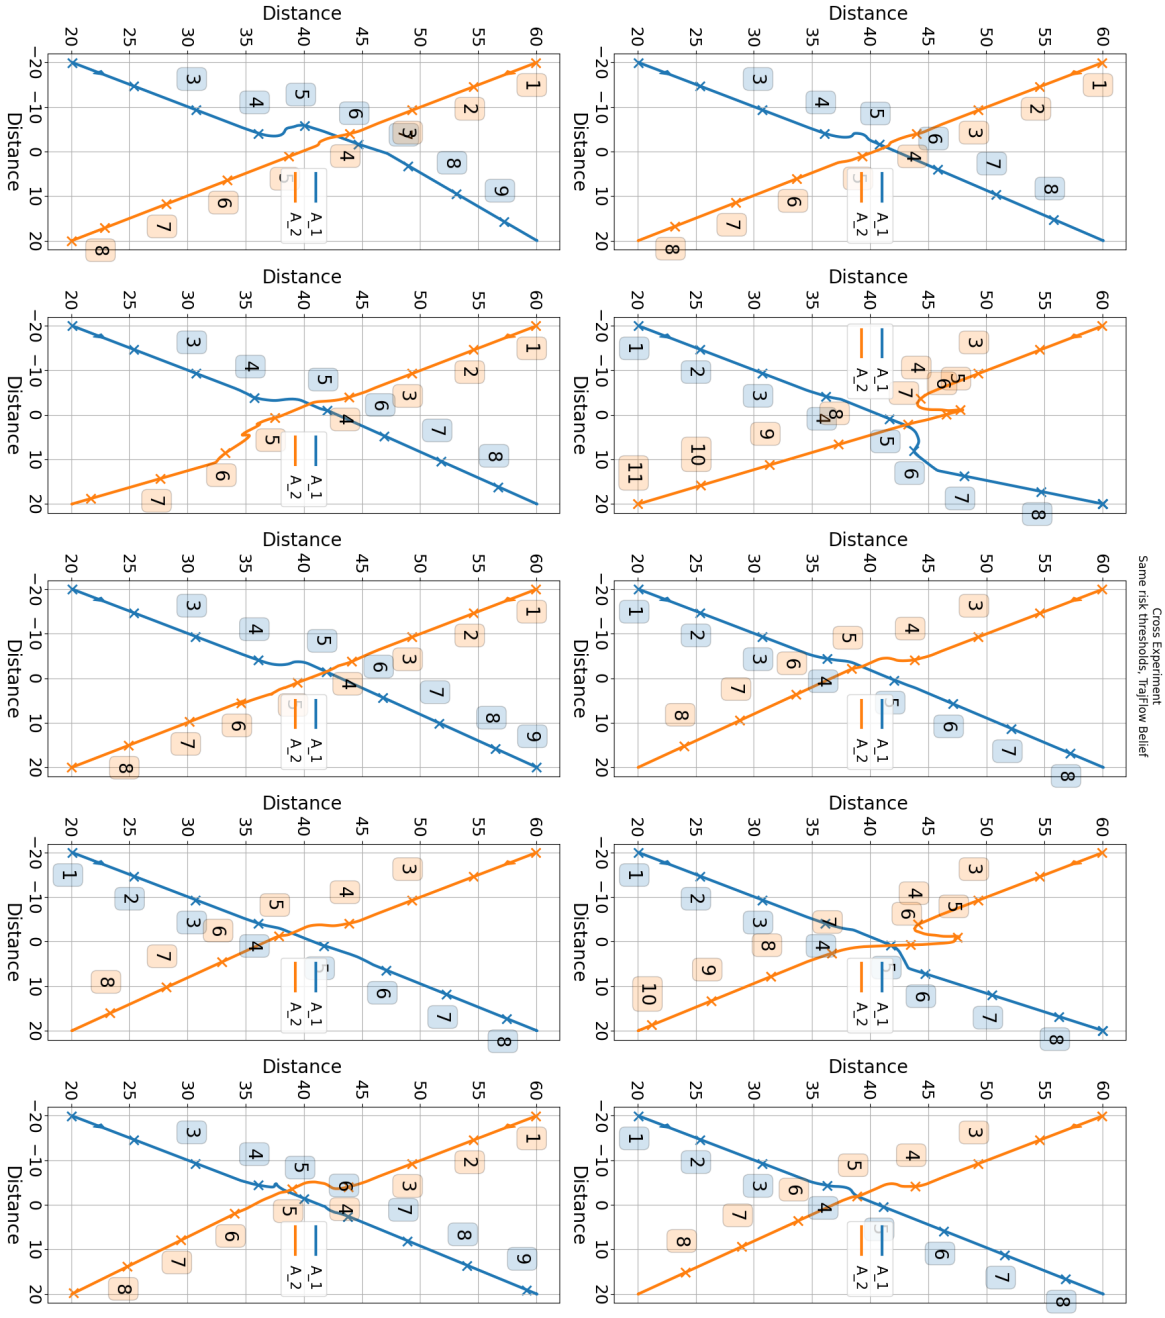


Figure B.9: Experiment: Cross, Risk threshold: Equal, Belief: TrajFlow

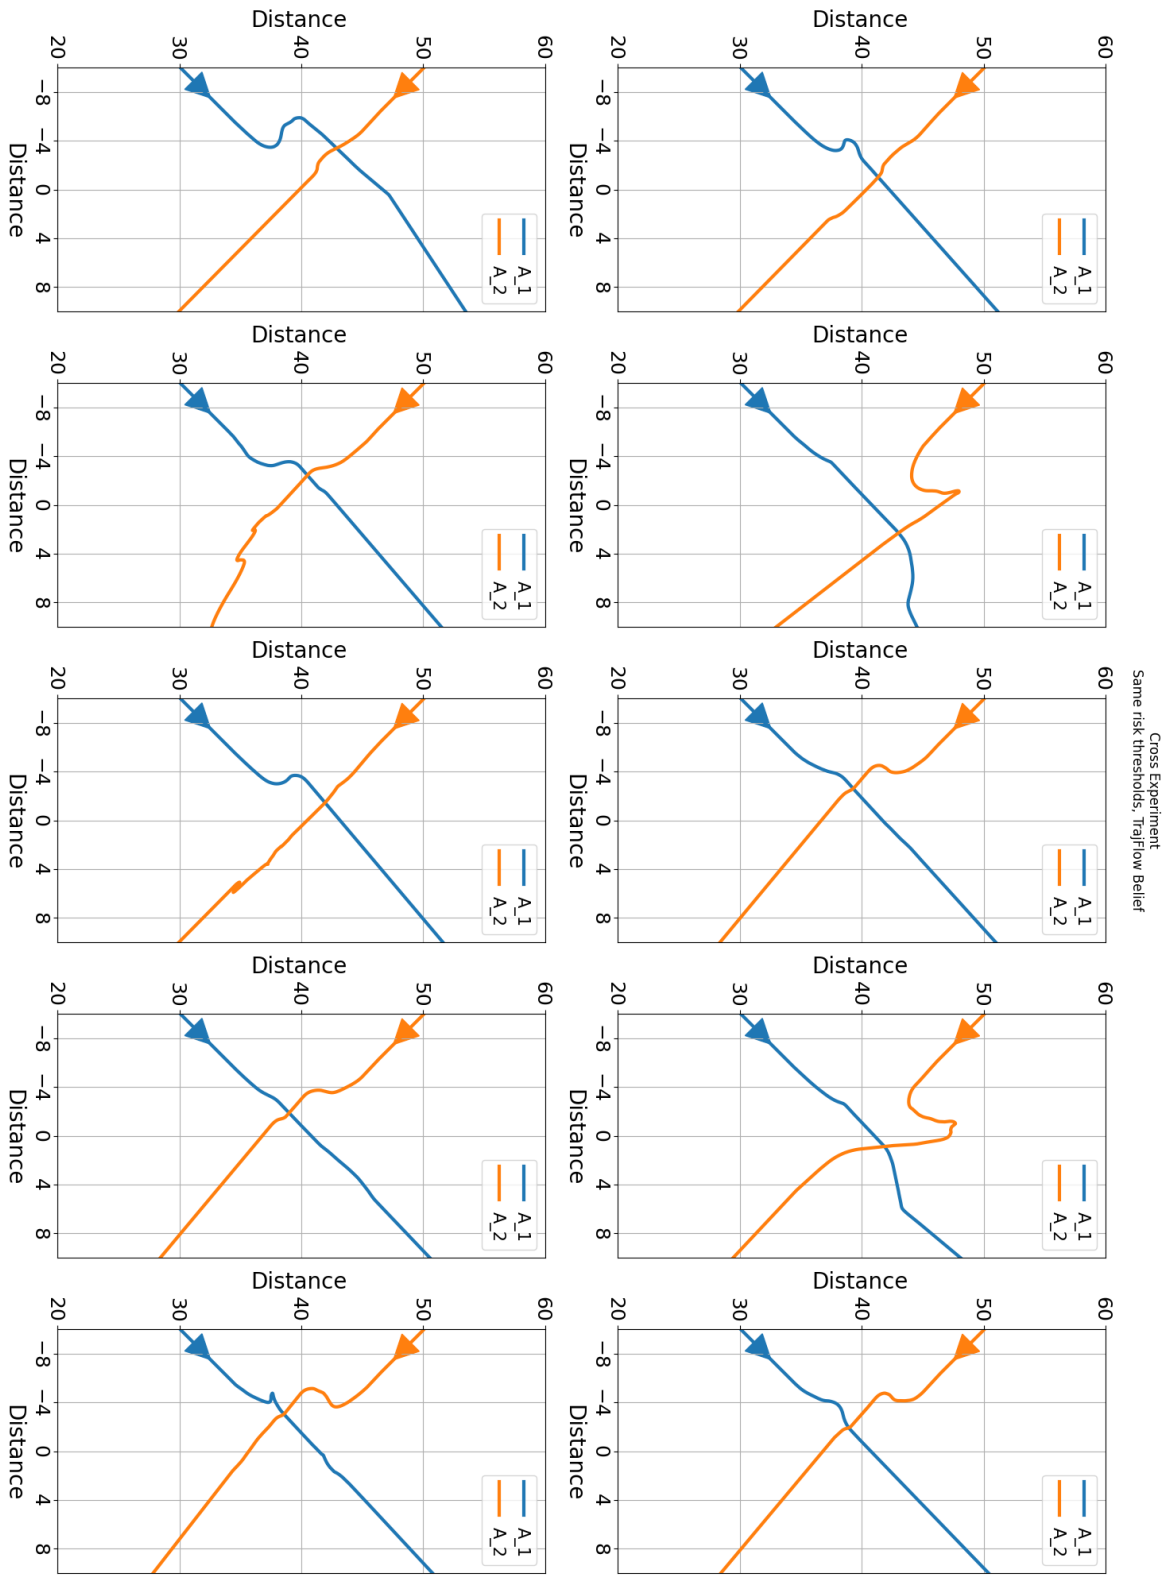


Figure B.10: Experiment: Cross, Risk threshold: Equal, Belief: TrajFlow, magnified

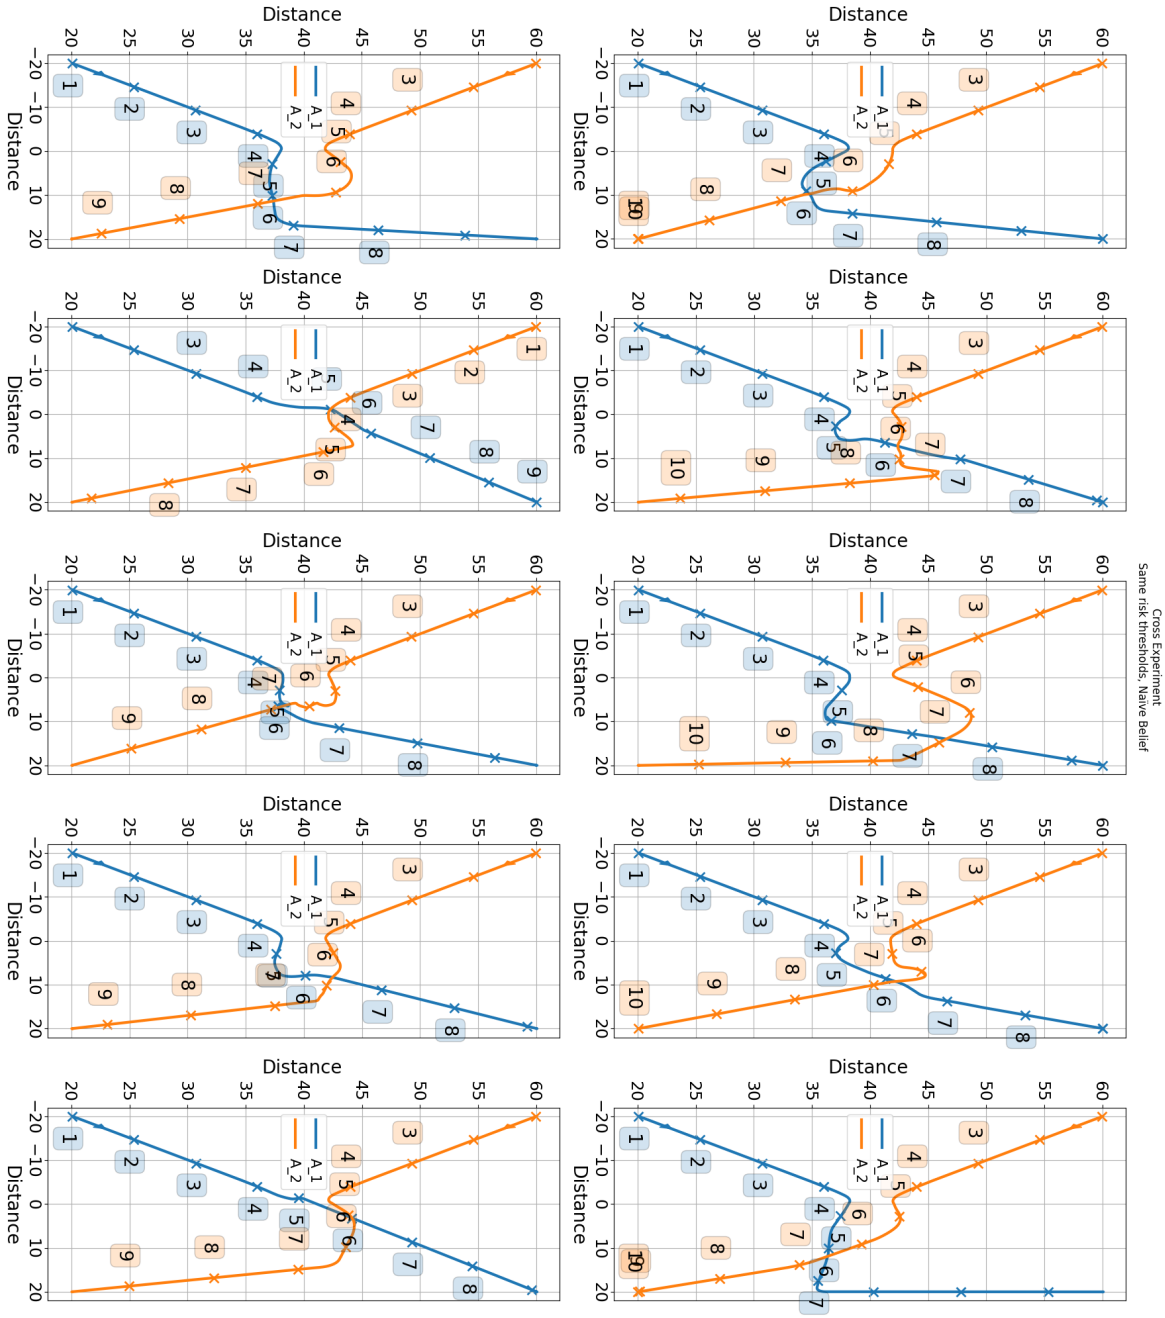


Figure B.11: Experiment: Cross, Risk threshold: Equal, Belief: Naïve

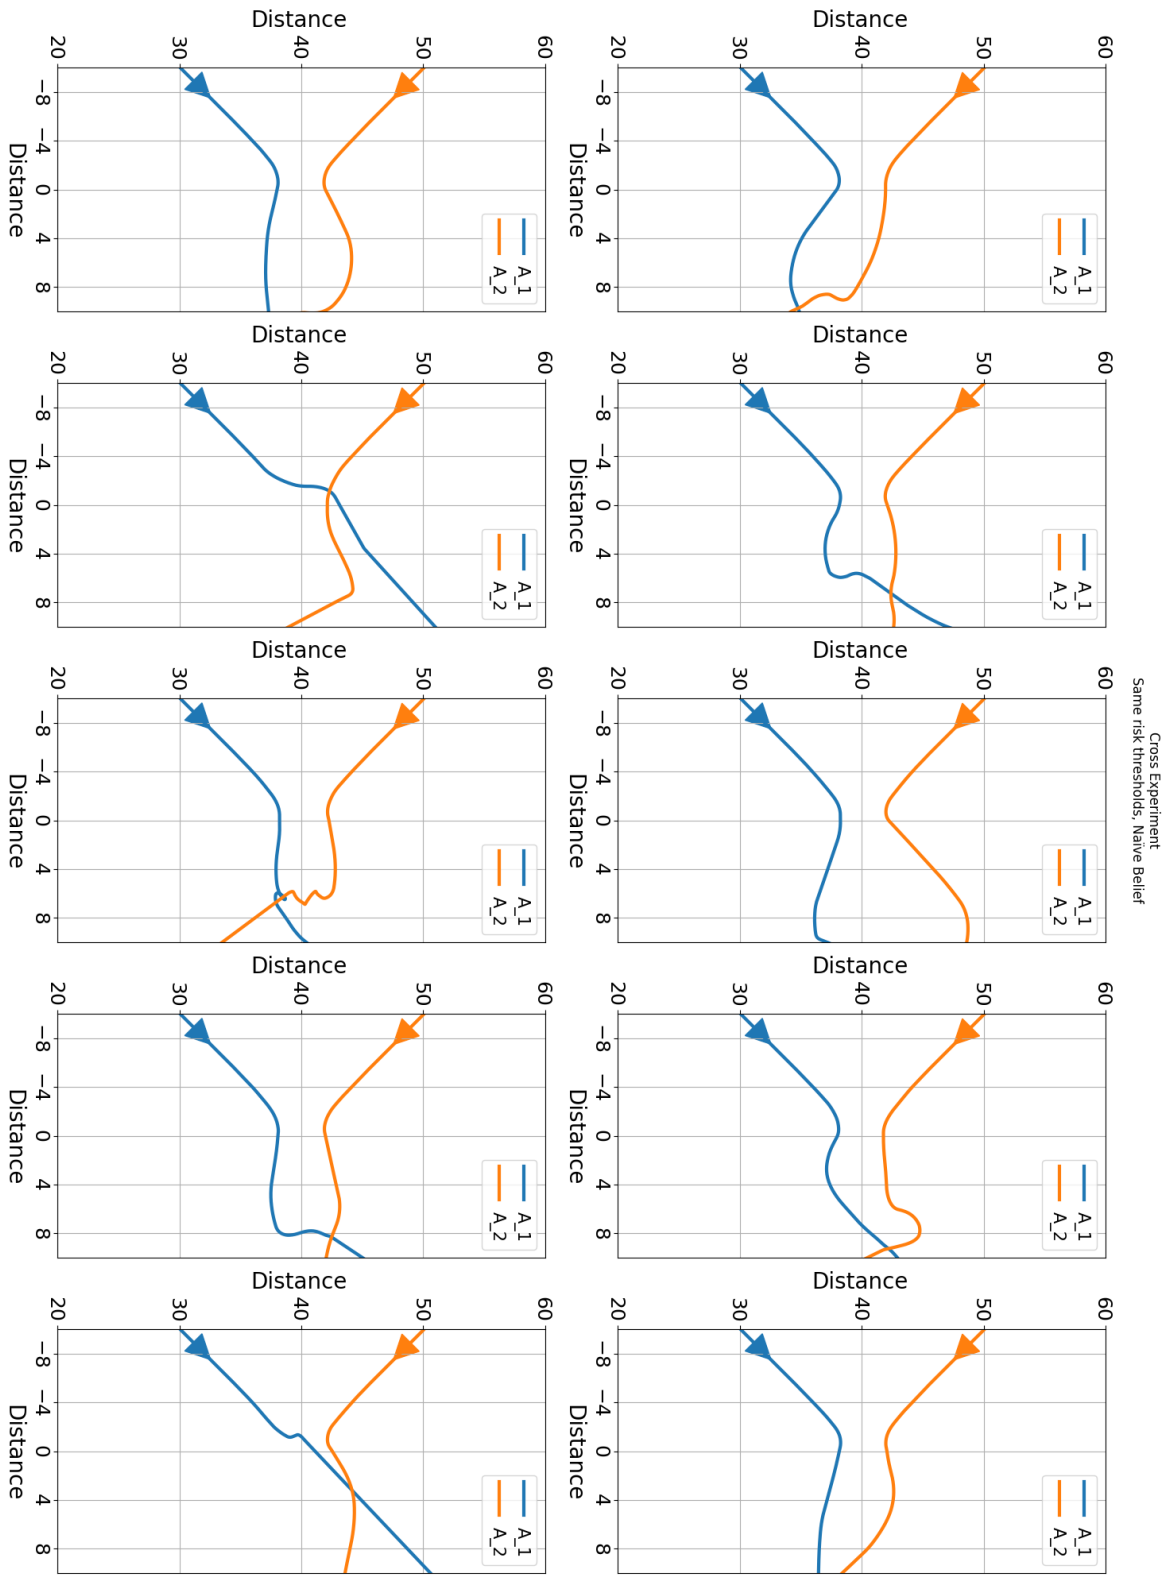


Figure B.12: Experiment: Cross, Risk threshold: Equal, Belief: Naïve, magnified

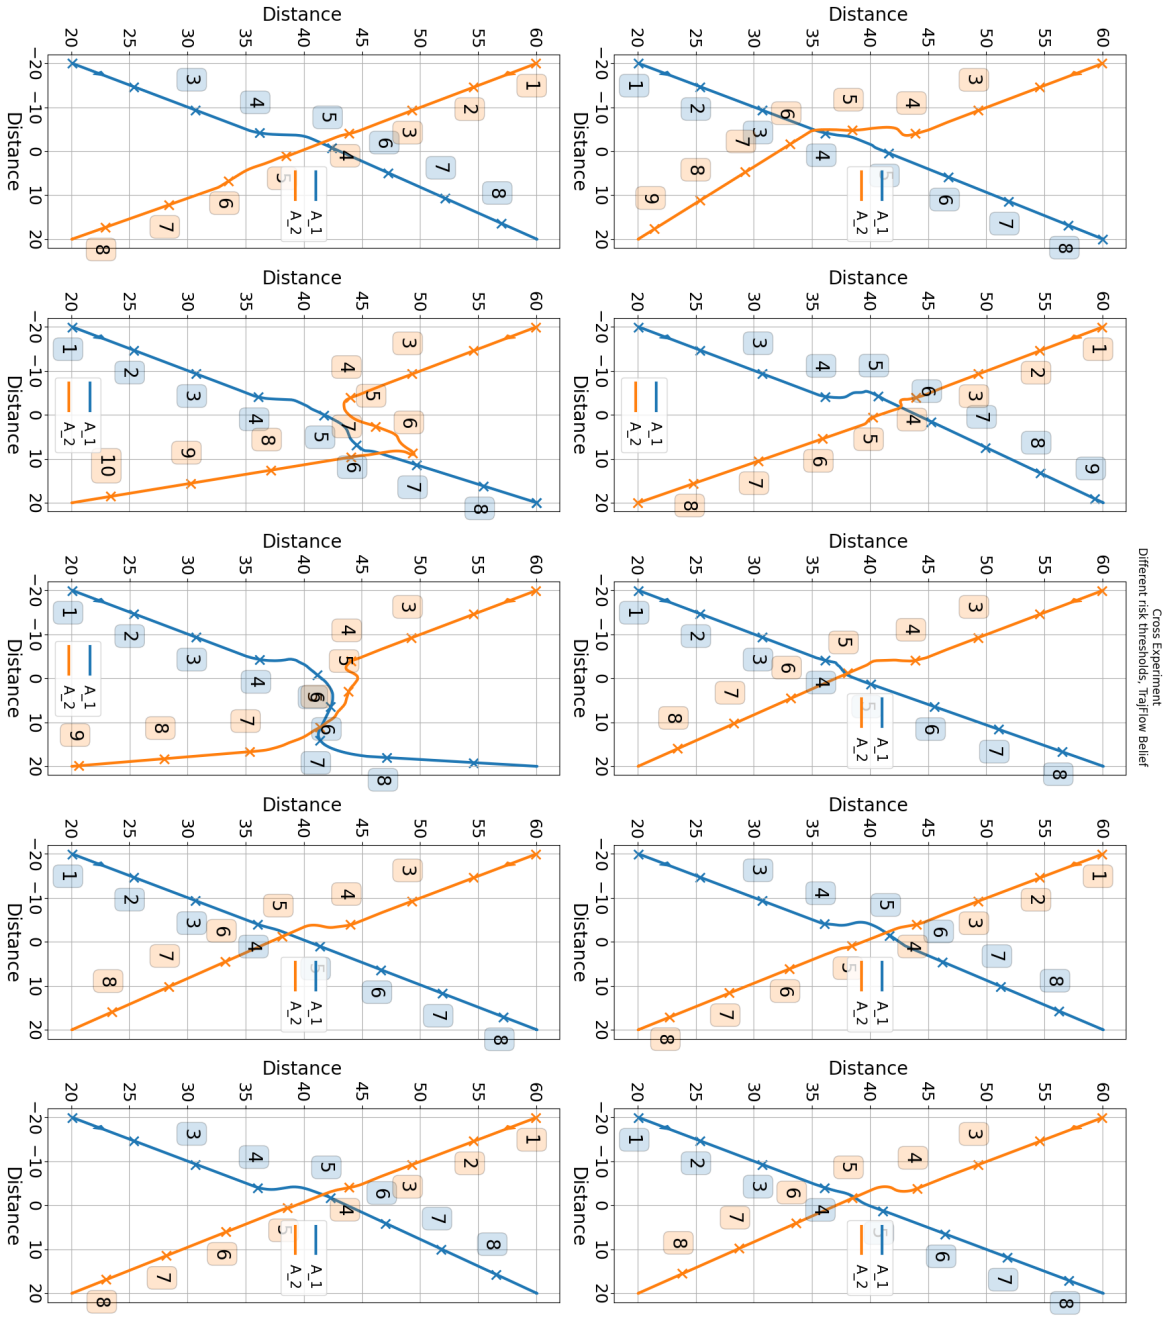


Figure B.13: Experiment: Cross, Risk threshold: Different, Belief: TrajFlow

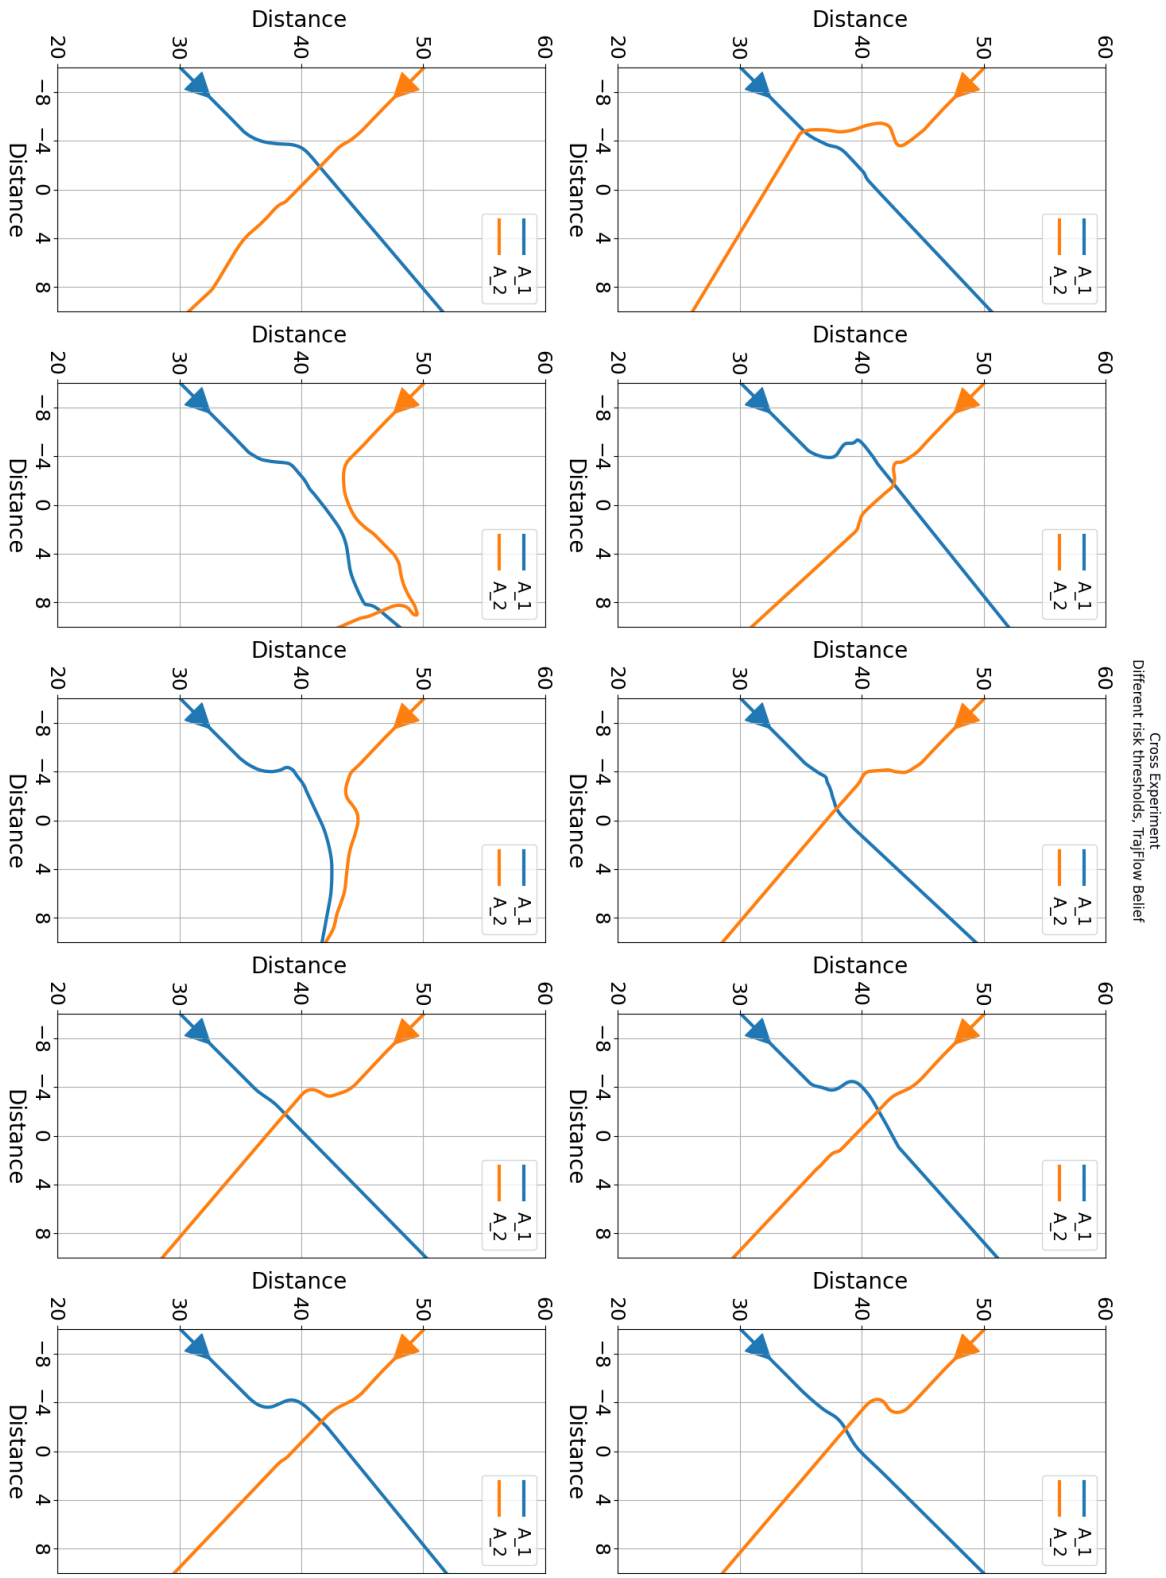


Figure B.14: Experiment: Cross, Risk threshold: Different, Belief: TrajFlow, magnified

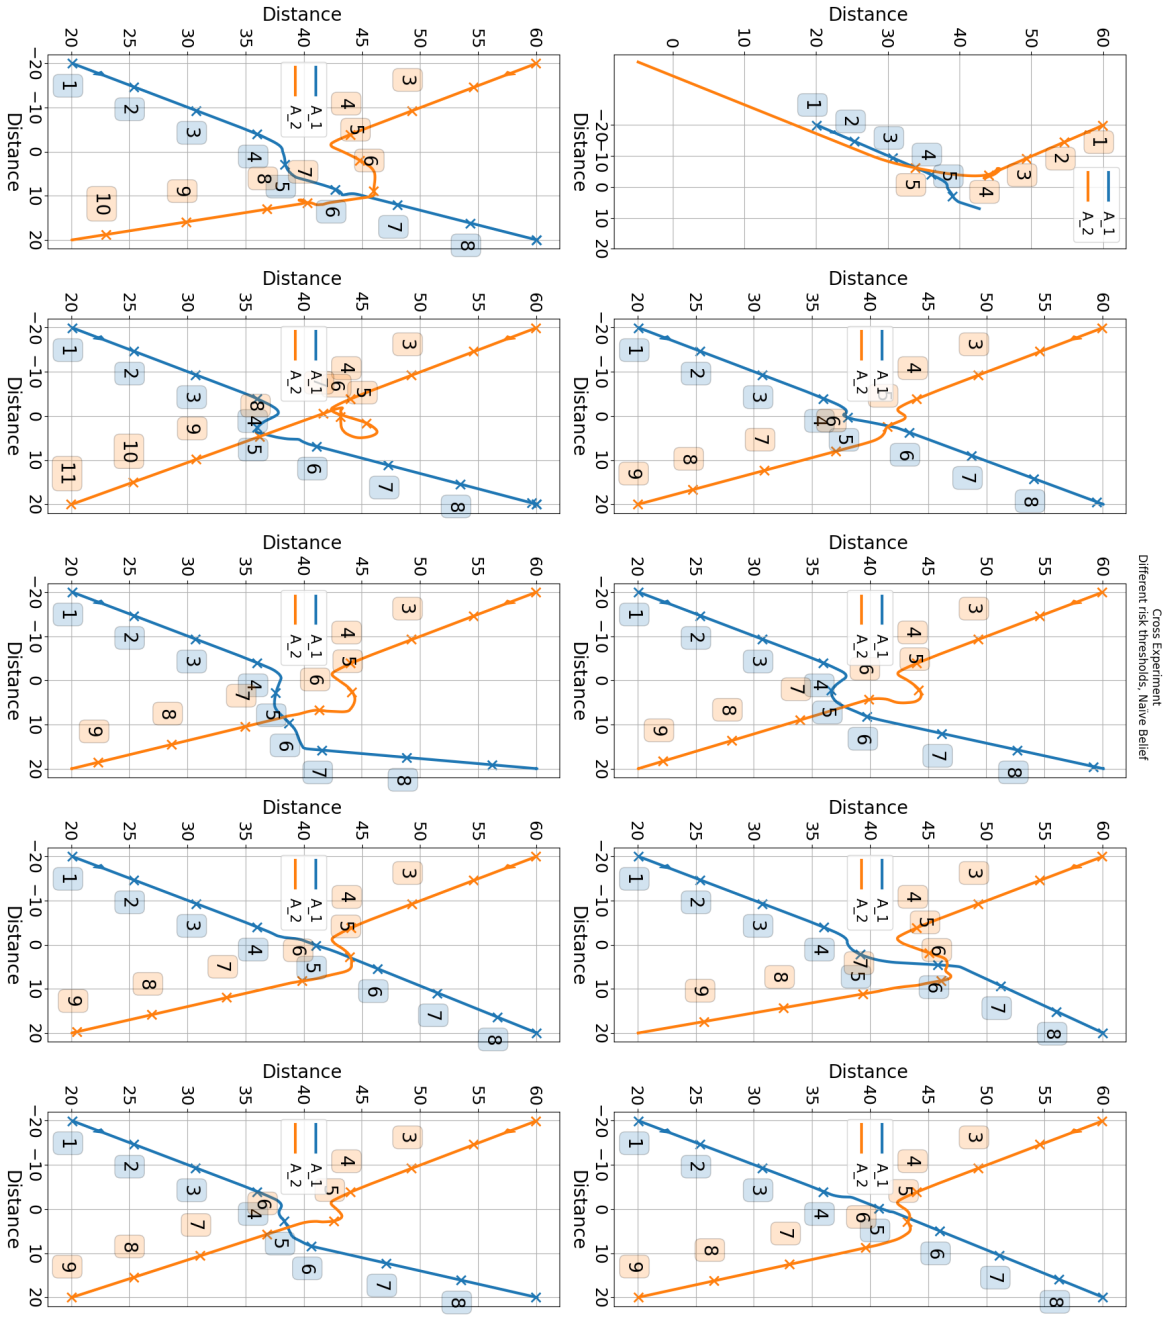


Figure B.15: Experiment: Cross, Risk threshold: Different, Belief: Naïve

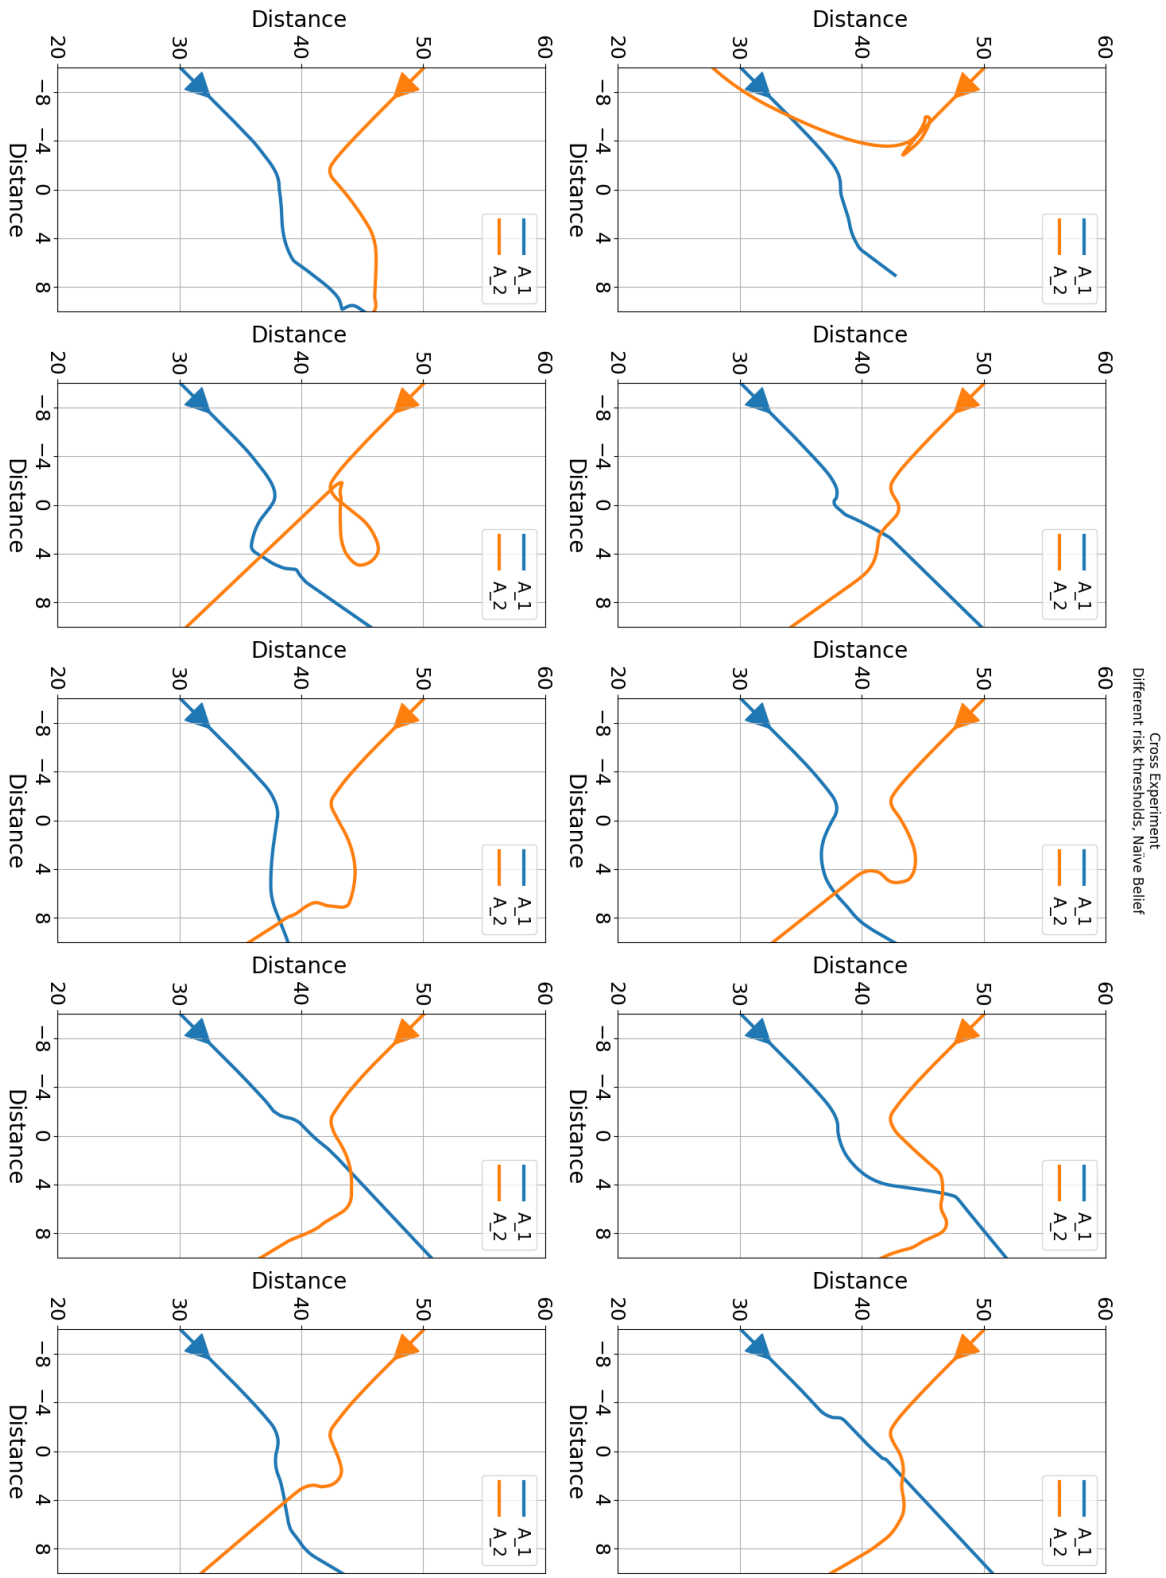


Figure B.16: Experiment: Cross, Risk threshold: Different, Belief: Naïve, magnified

## B.3. Experiment: Overtake

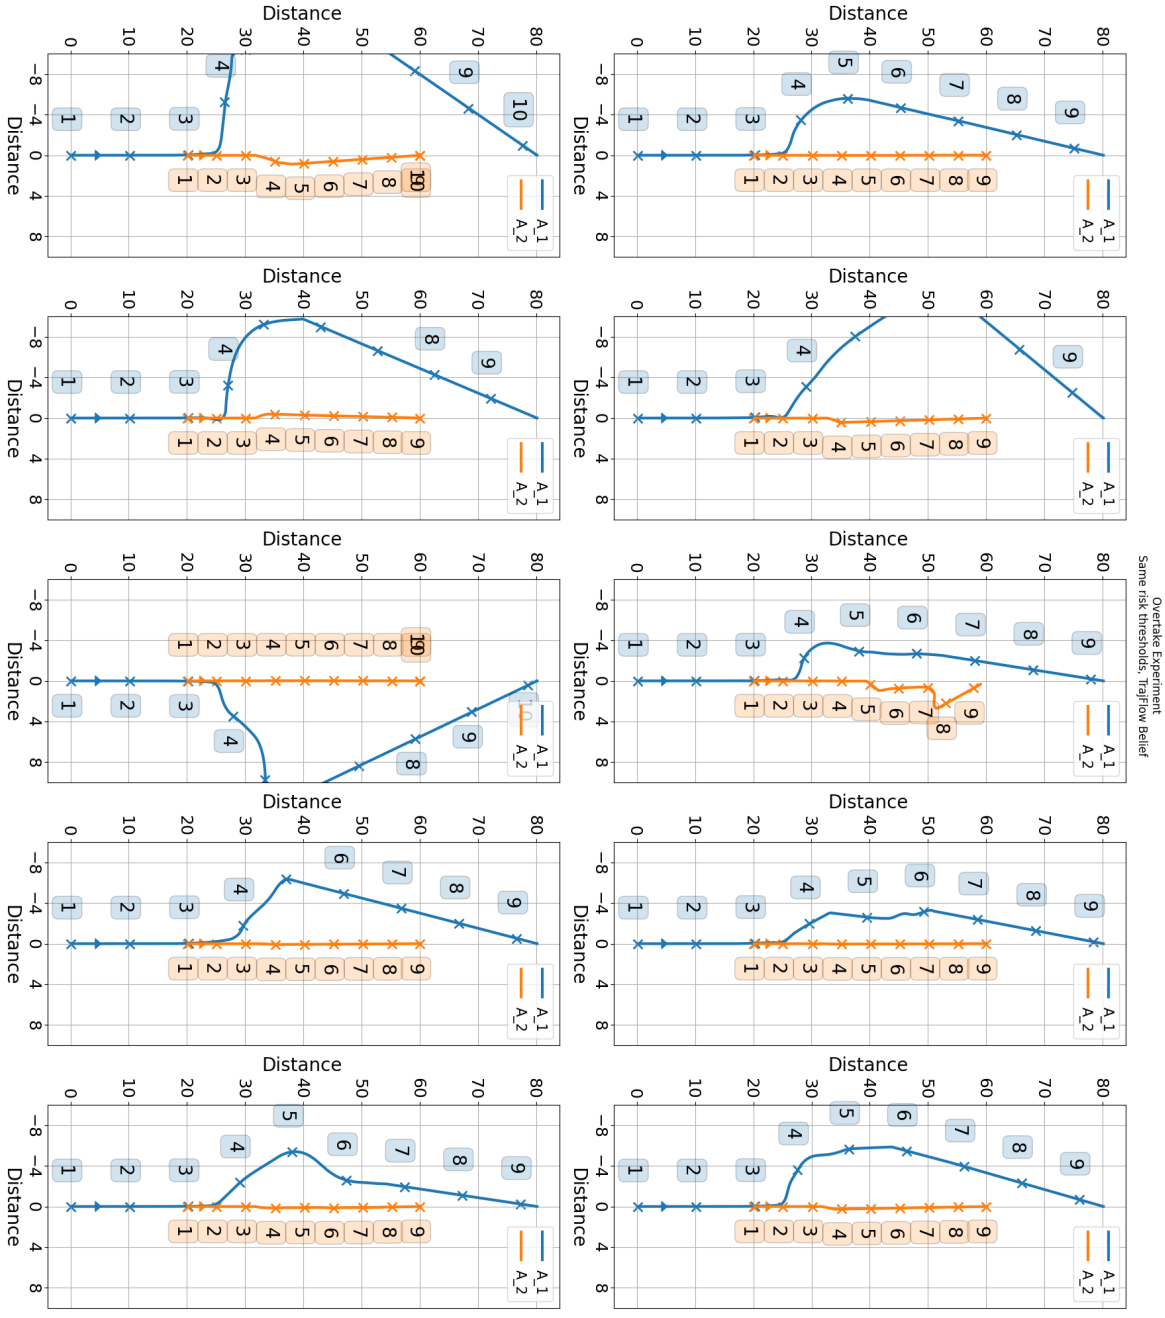


Figure B.17: Experiment: Overtake, Risk threshold: Equal, Belief: TrajFlow

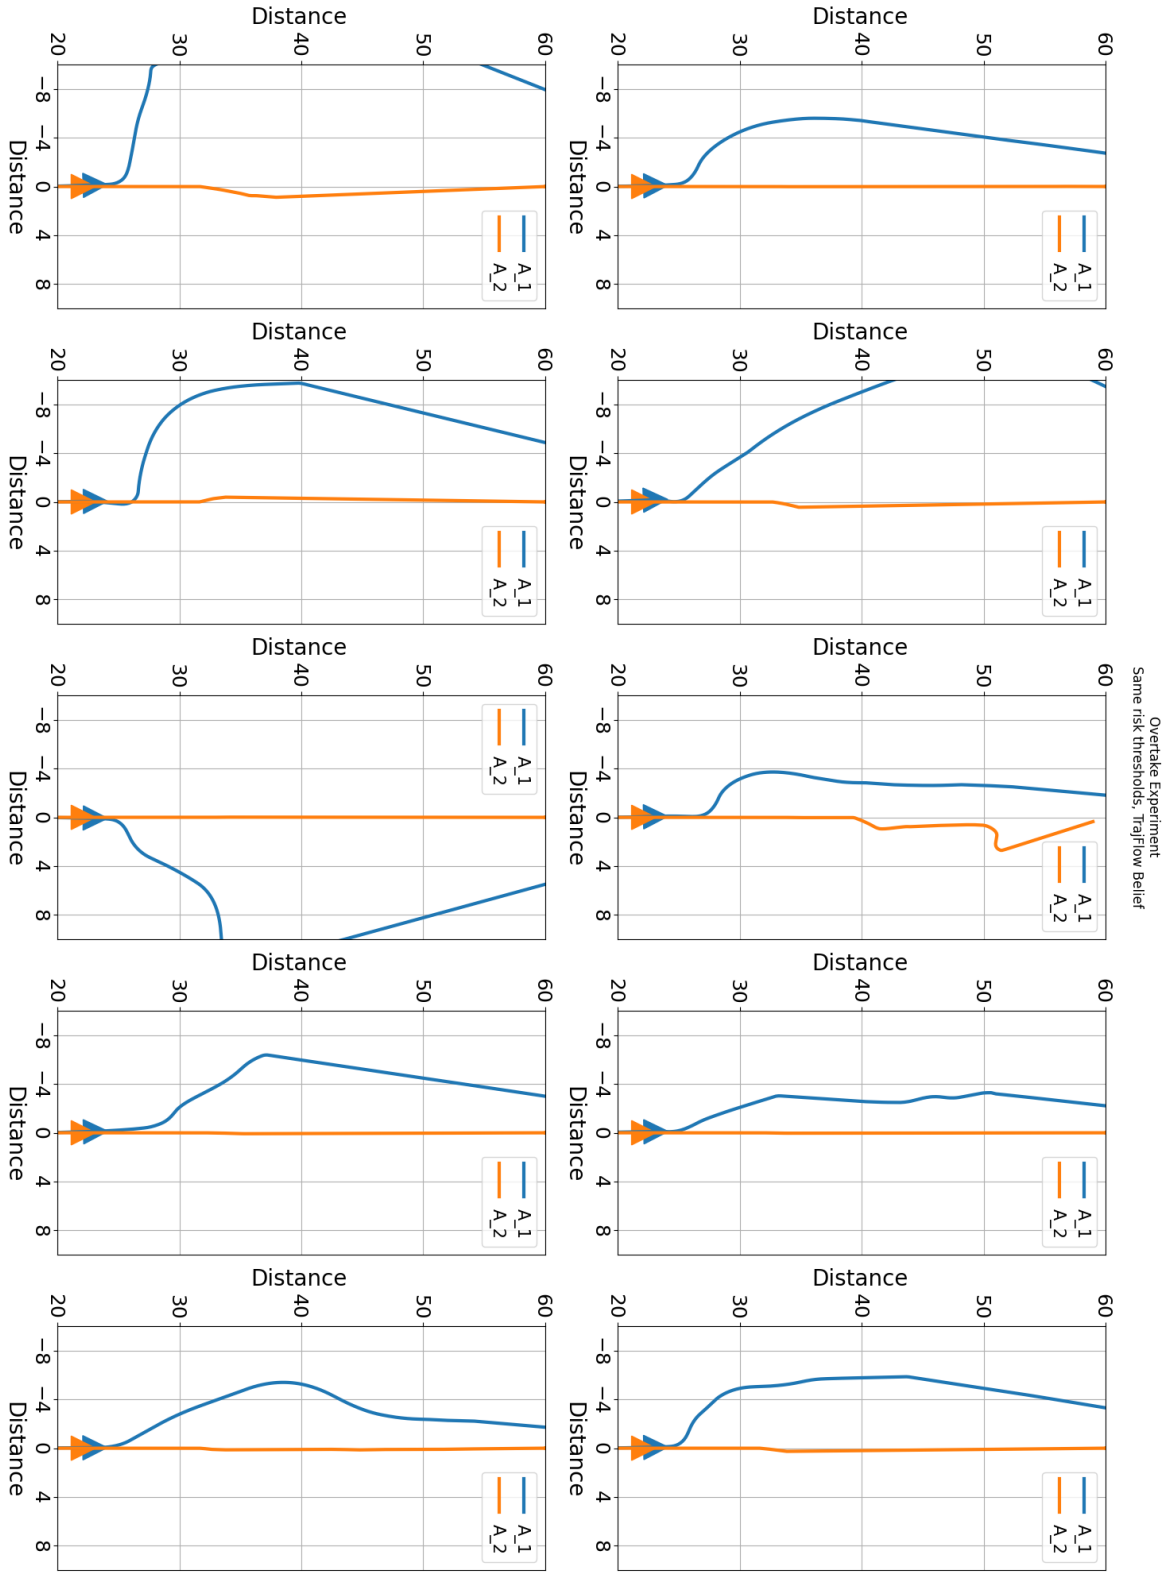


Figure B.18: Experiment: Overtake, Risk threshold: Equal, Belief: TrajFlow, magnified

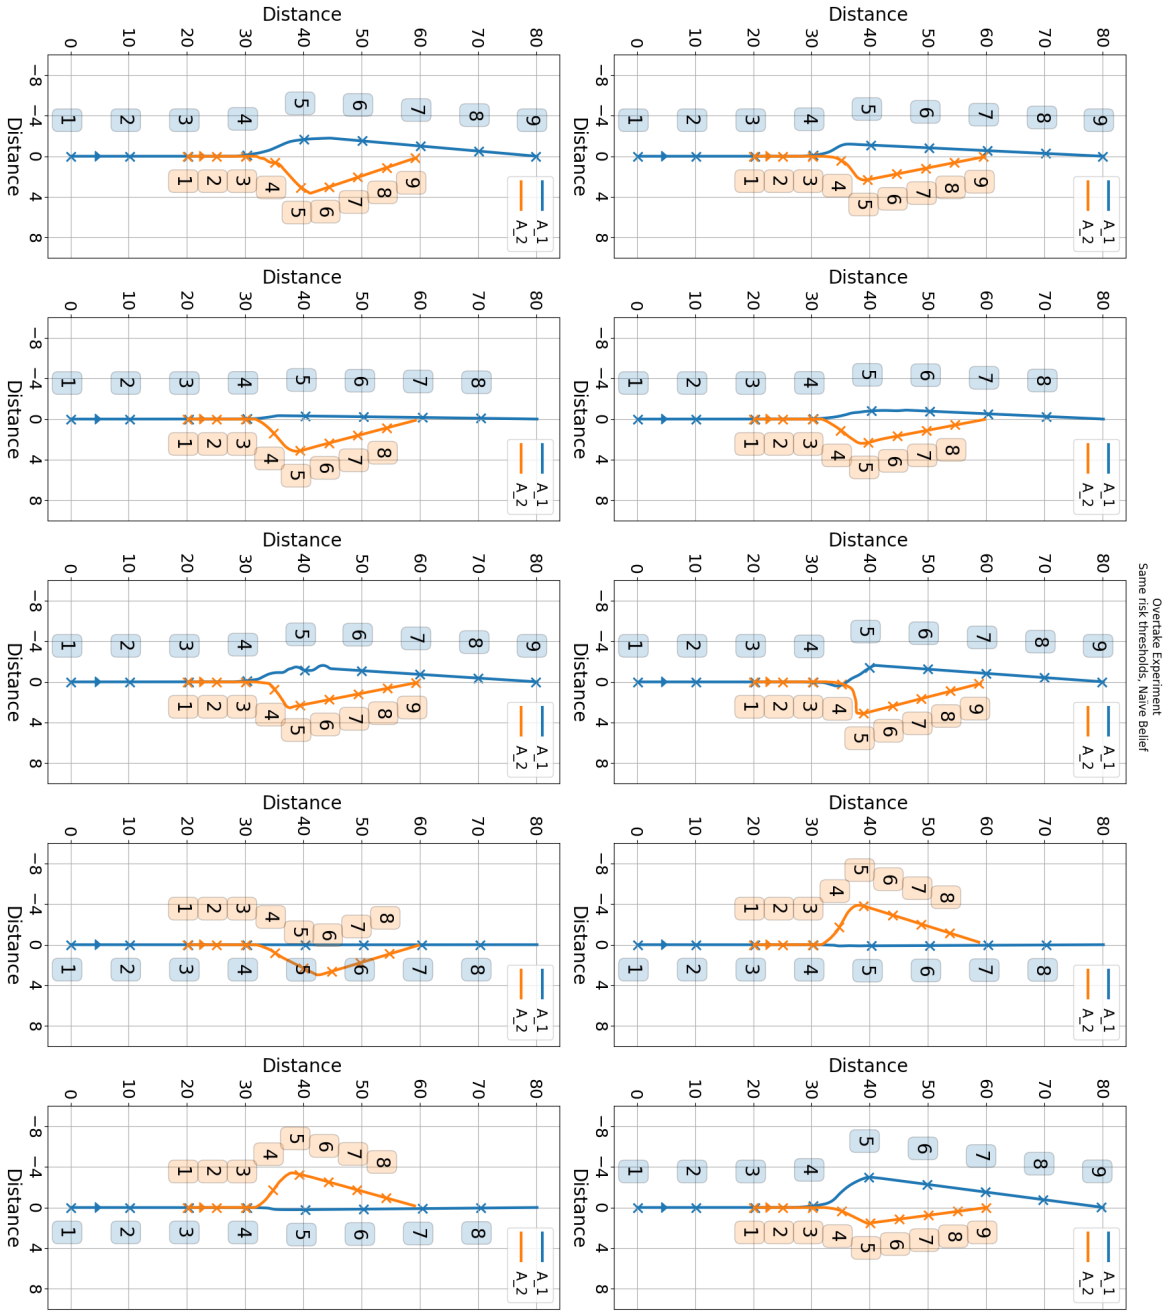


Figure B.19: Experiment: Overtake, Risk threshold: Equal, Belief: Naïve

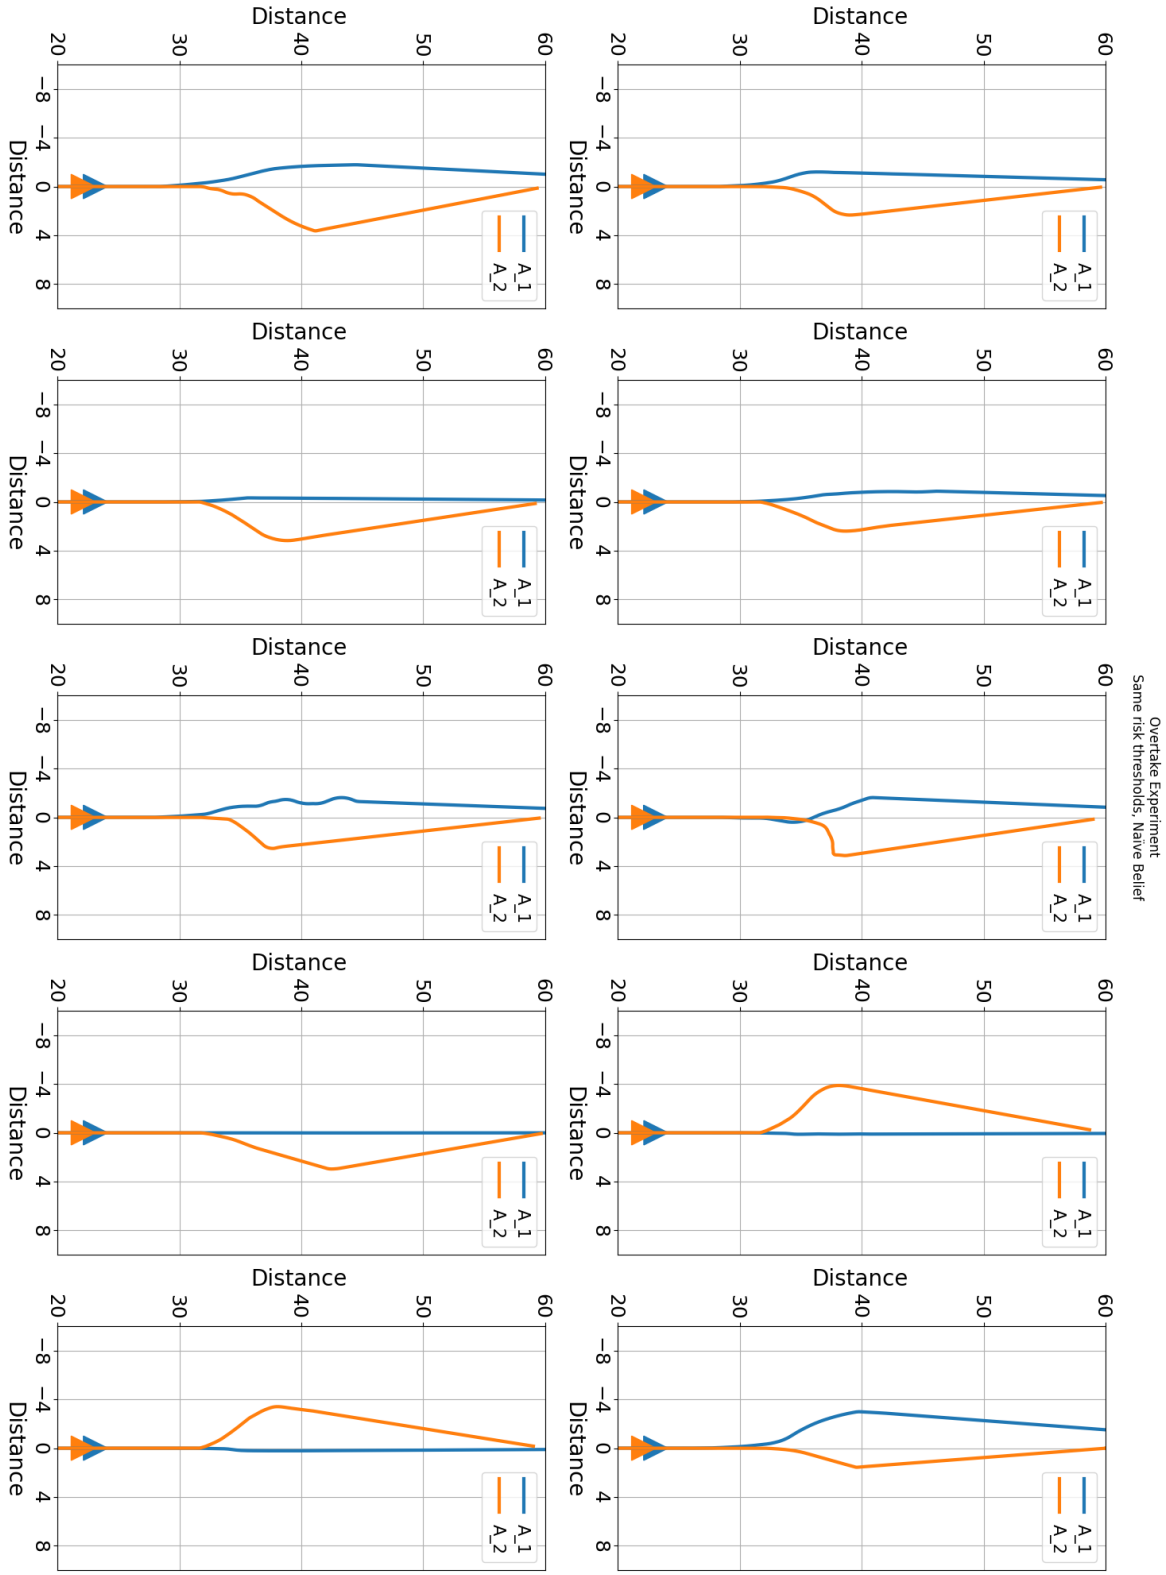


Figure B.20: Experiment: Overtake, Risk threshold: Equal, Belief: Naïve, magnified

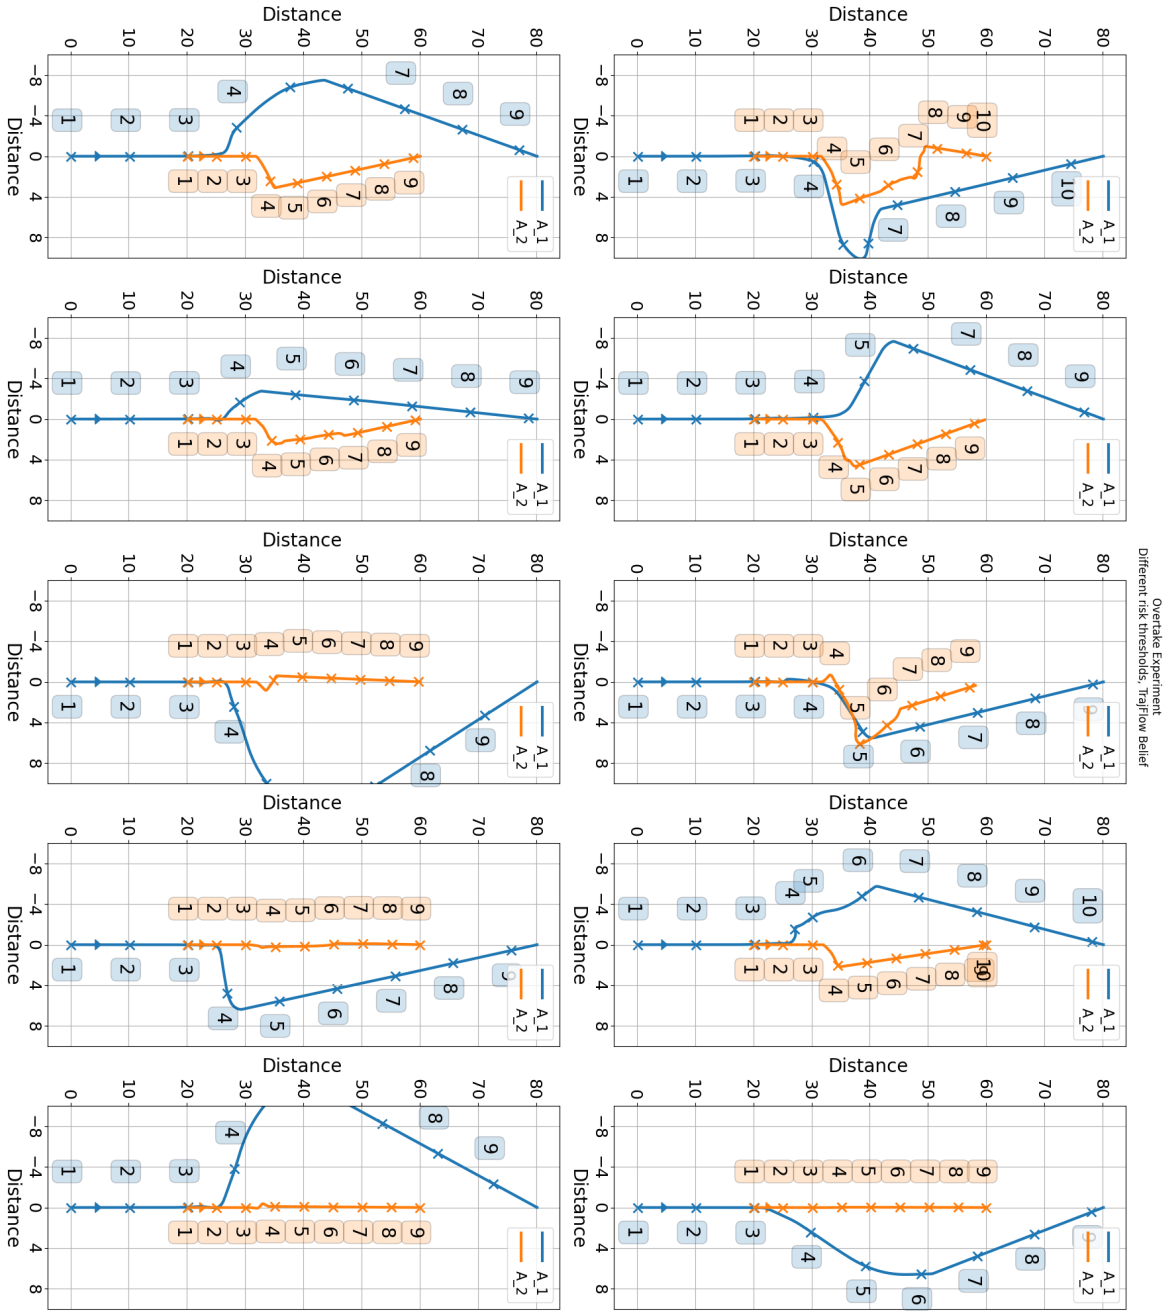


Figure B.21: Experiment: Overtake, Risk threshold: Different, Belief: TrajFlow

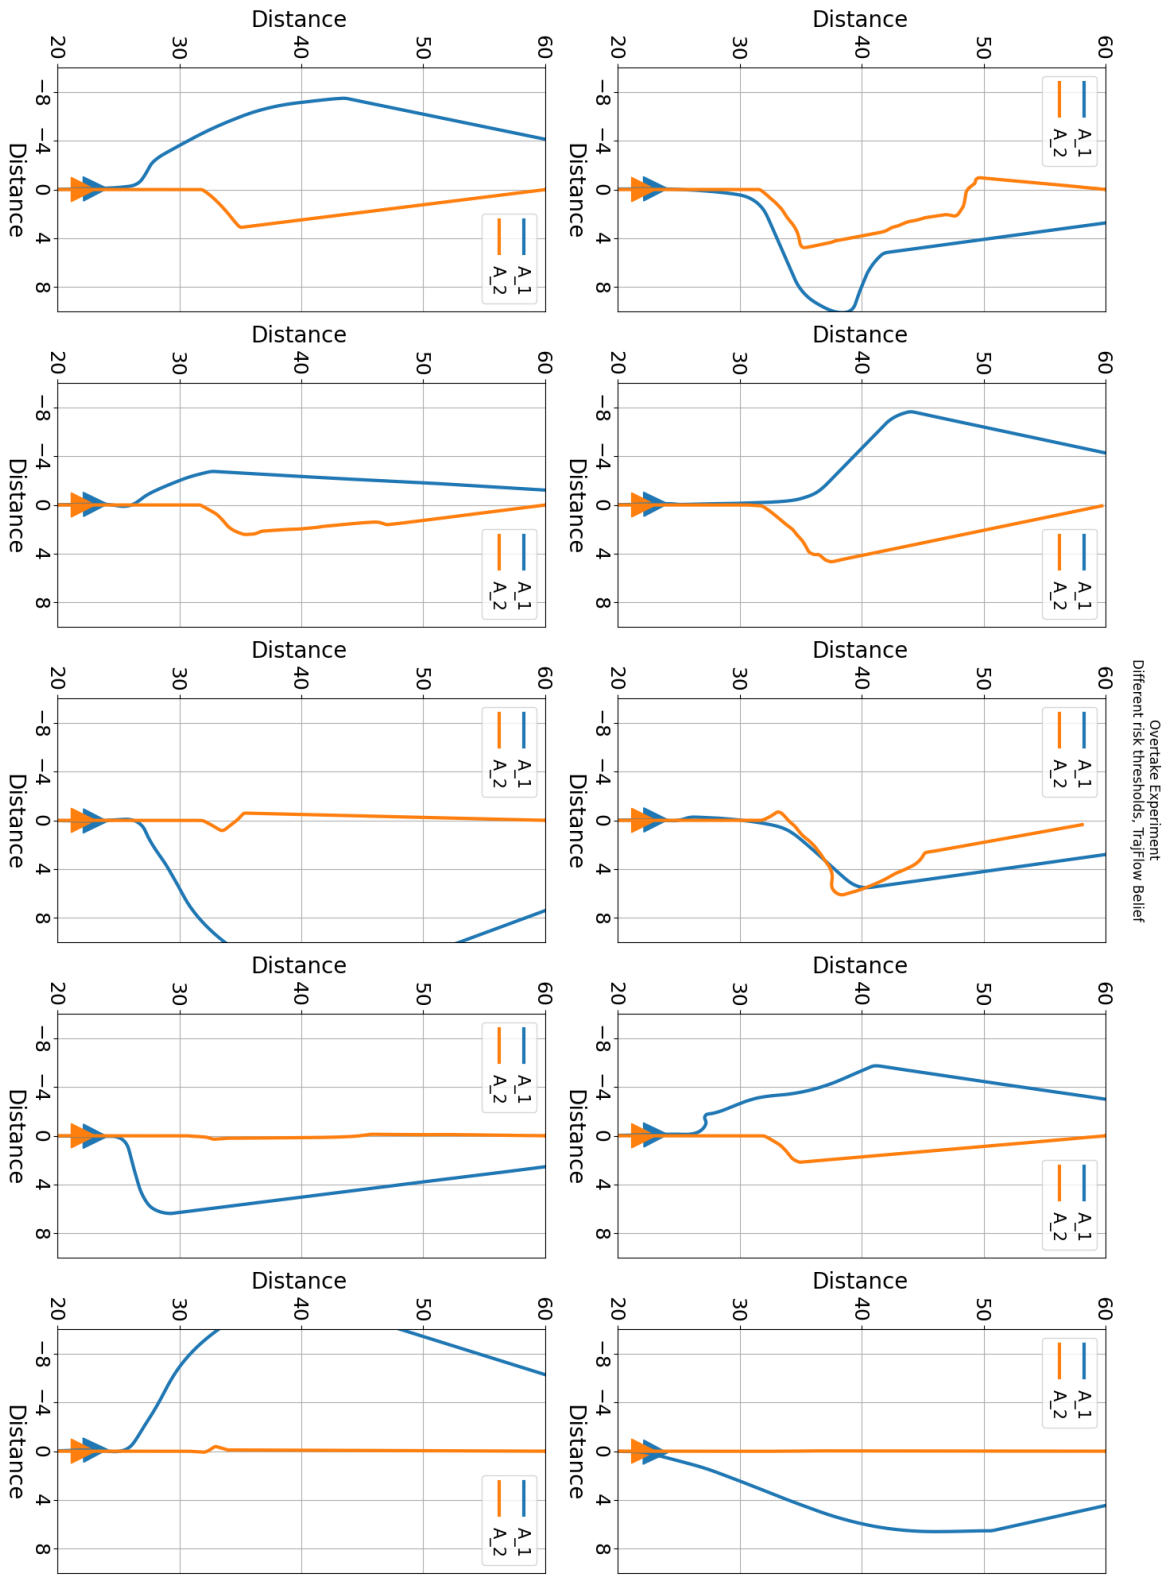


Figure B.22: Experiment: Overtake, Risk threshold: Different, Belief: TrajFlow, magnified

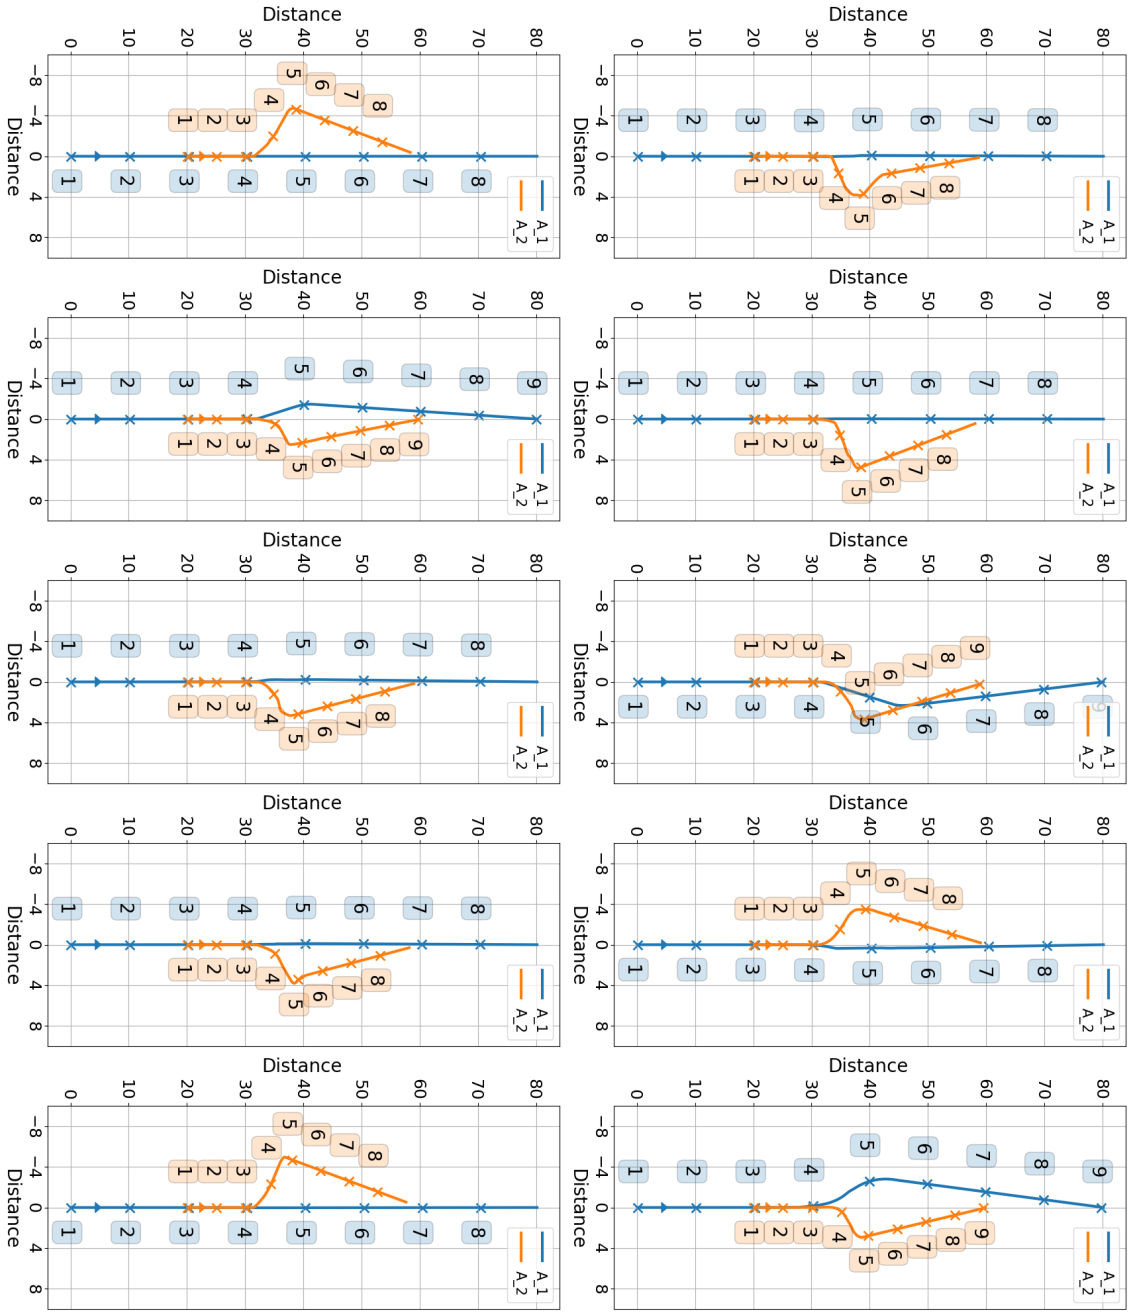


Figure B.23: Experiment: Overtake, Risk threshold: Different, Belief: Naïve

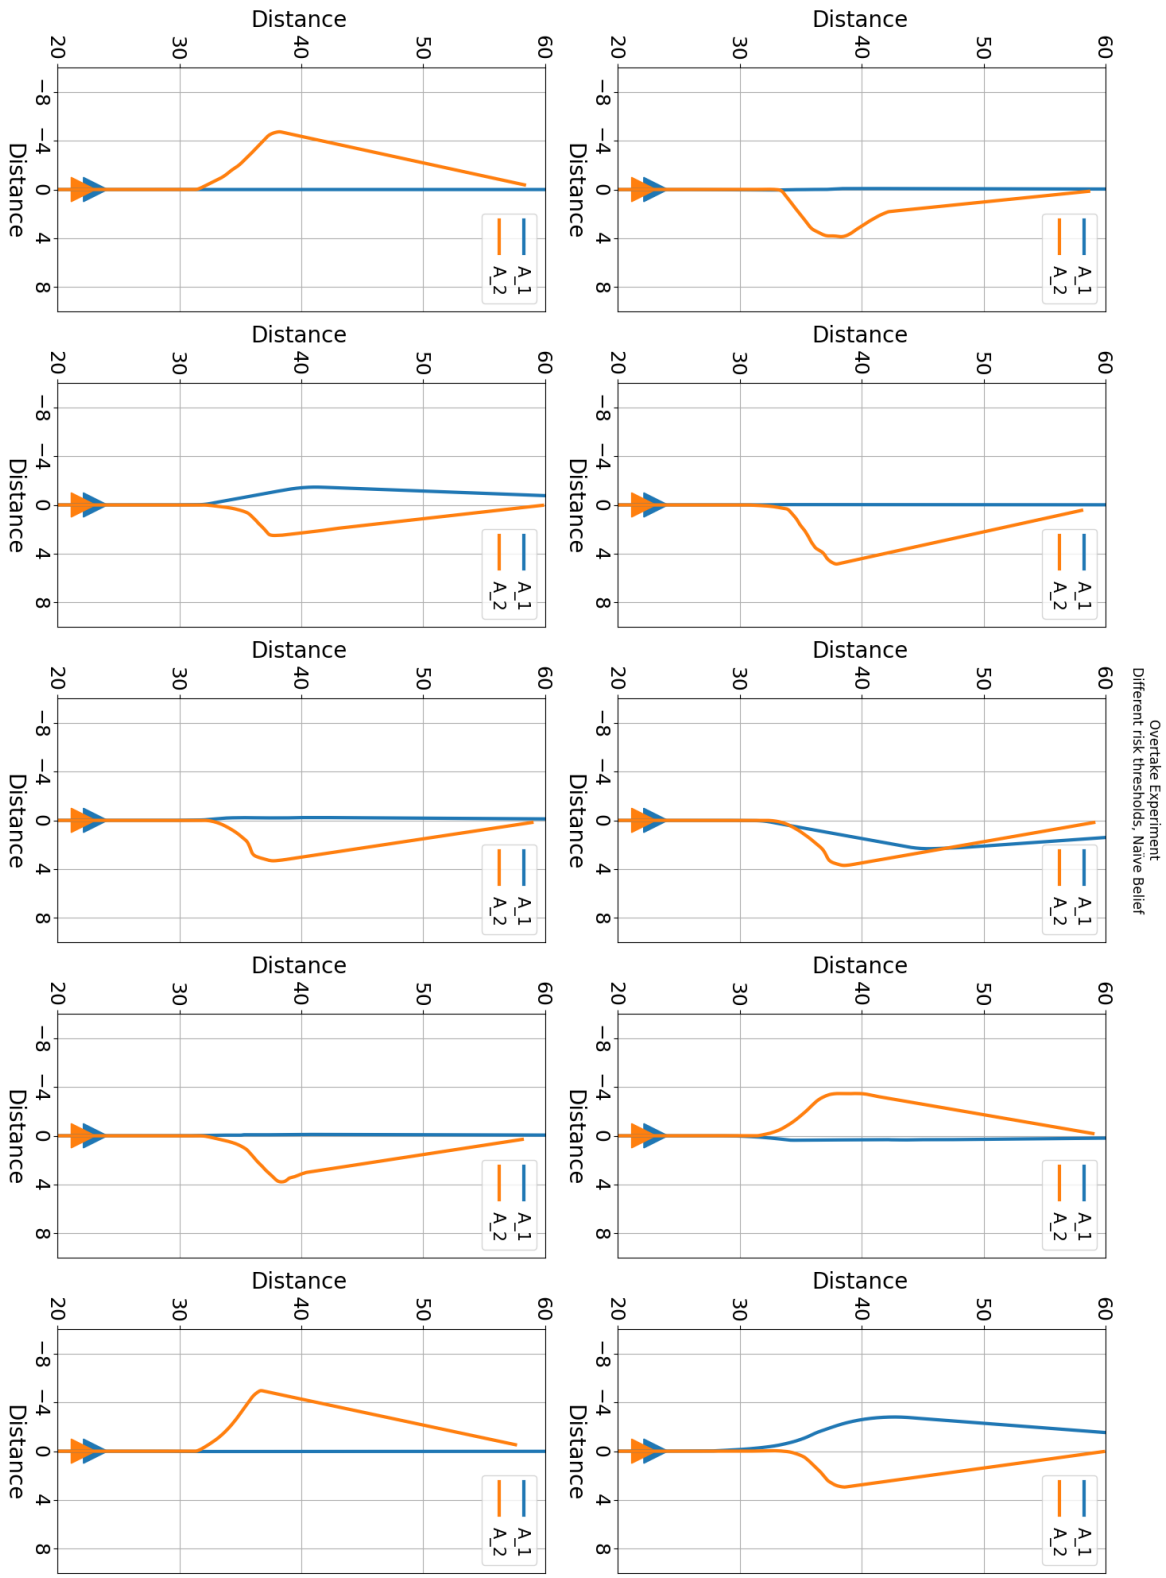


Figure B.24: Experiment: Overtake, Risk threshold: Different, Belief: Naïve, magnified

## B.4. Experiment: Same Goal

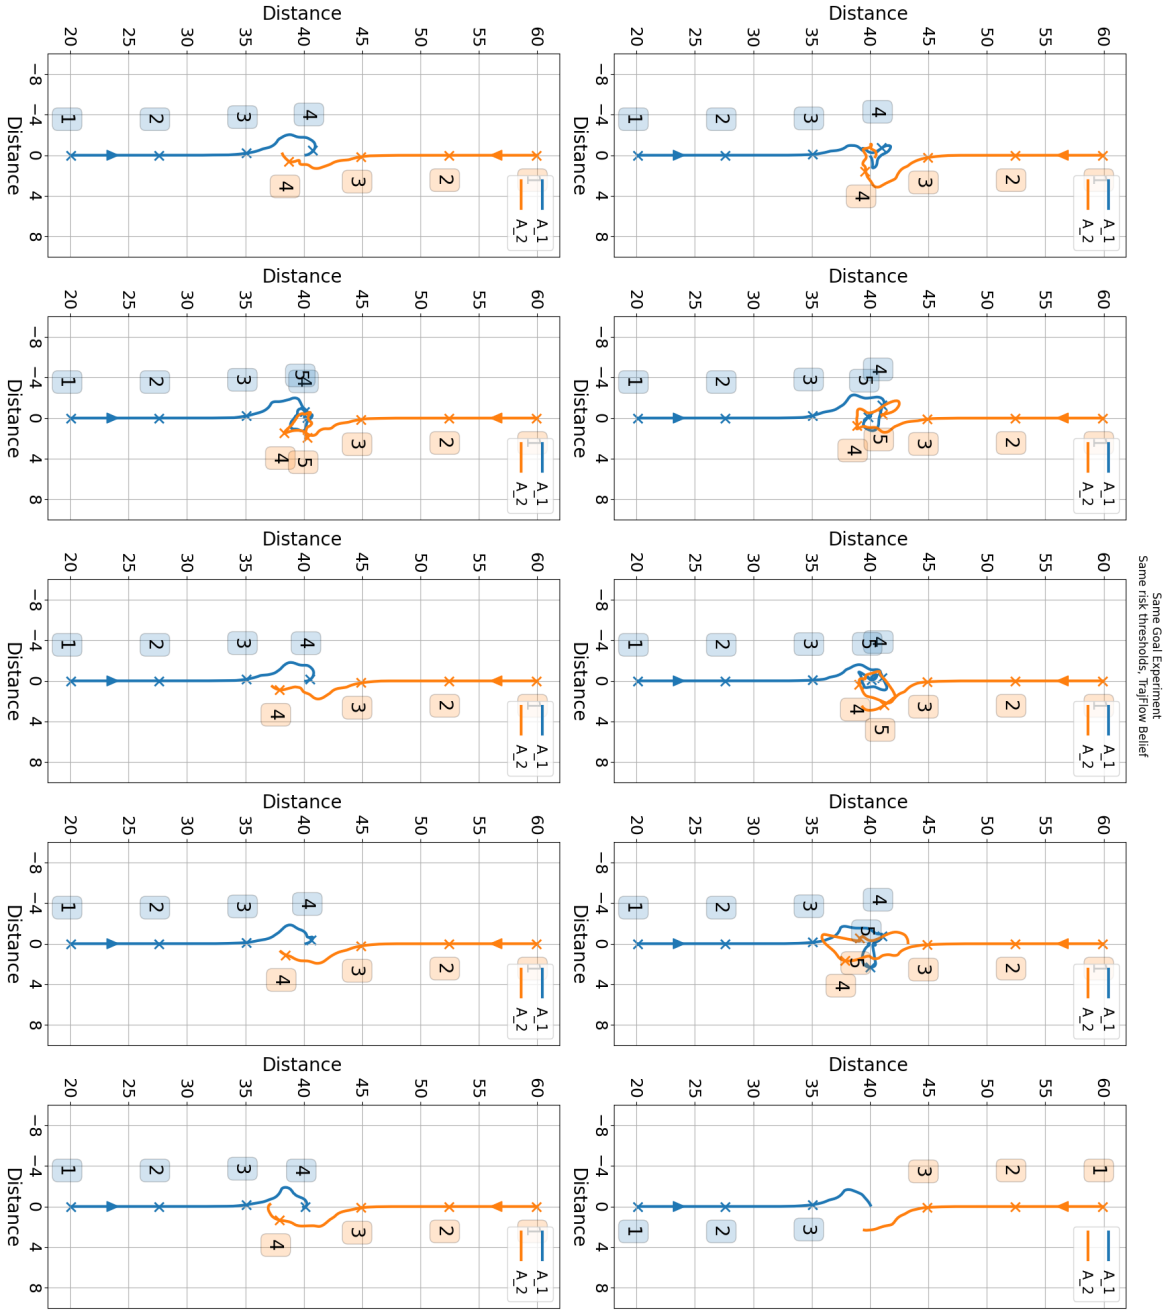


Figure B.25: Experiment: Same Goal, Risk threshold: Equal, Belief: TrajFlow

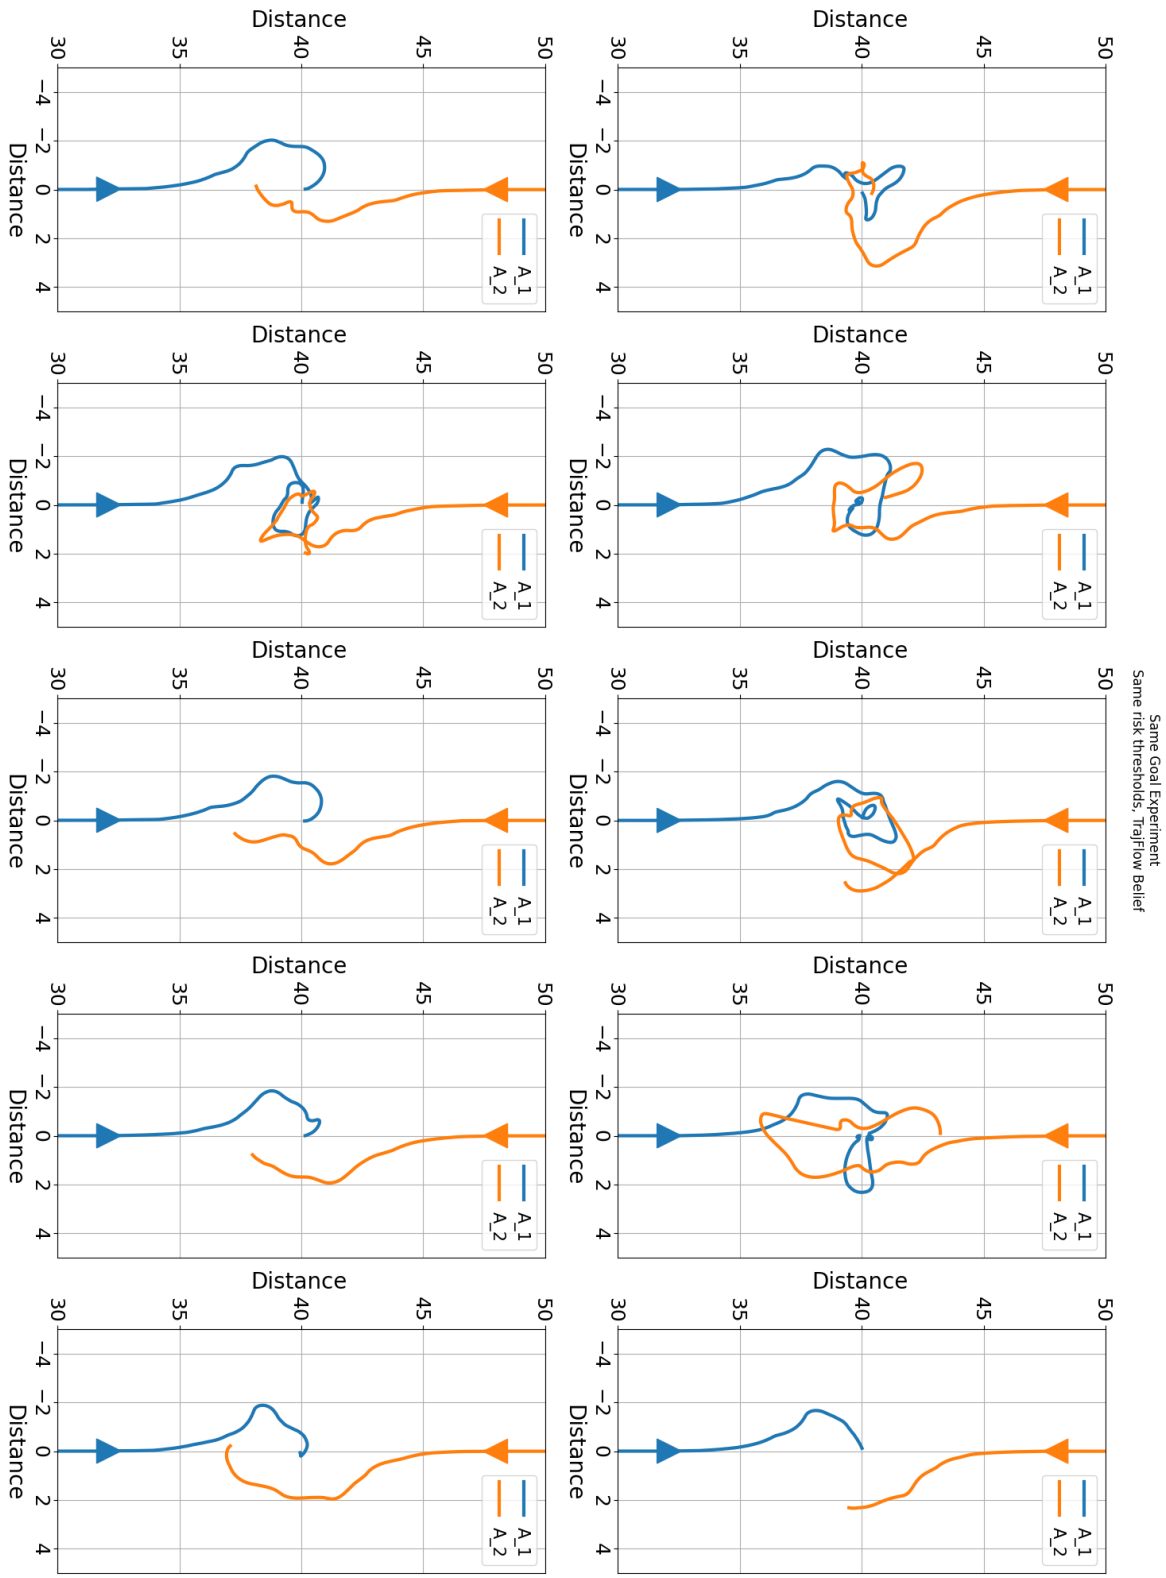


Figure B.26: Experiment: Same Goal, Risk threshold: Equal, Belief: TrajFlow, magnified

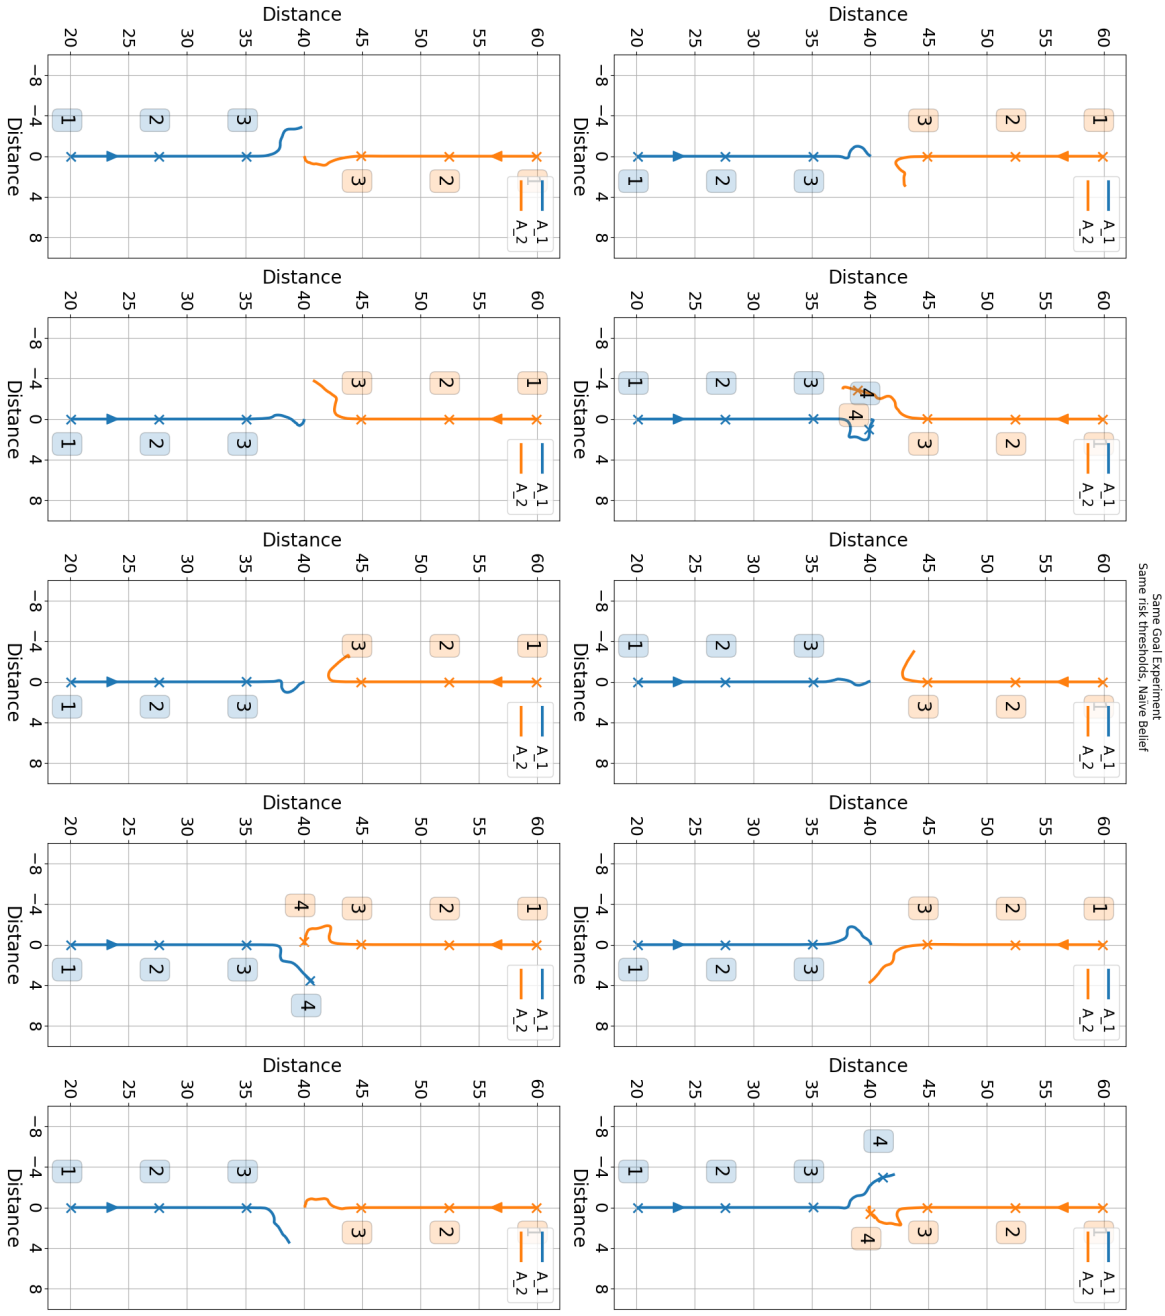


Figure B.27: Experiment: Same Goal, Risk threshold: Equal, Belief: Naïve

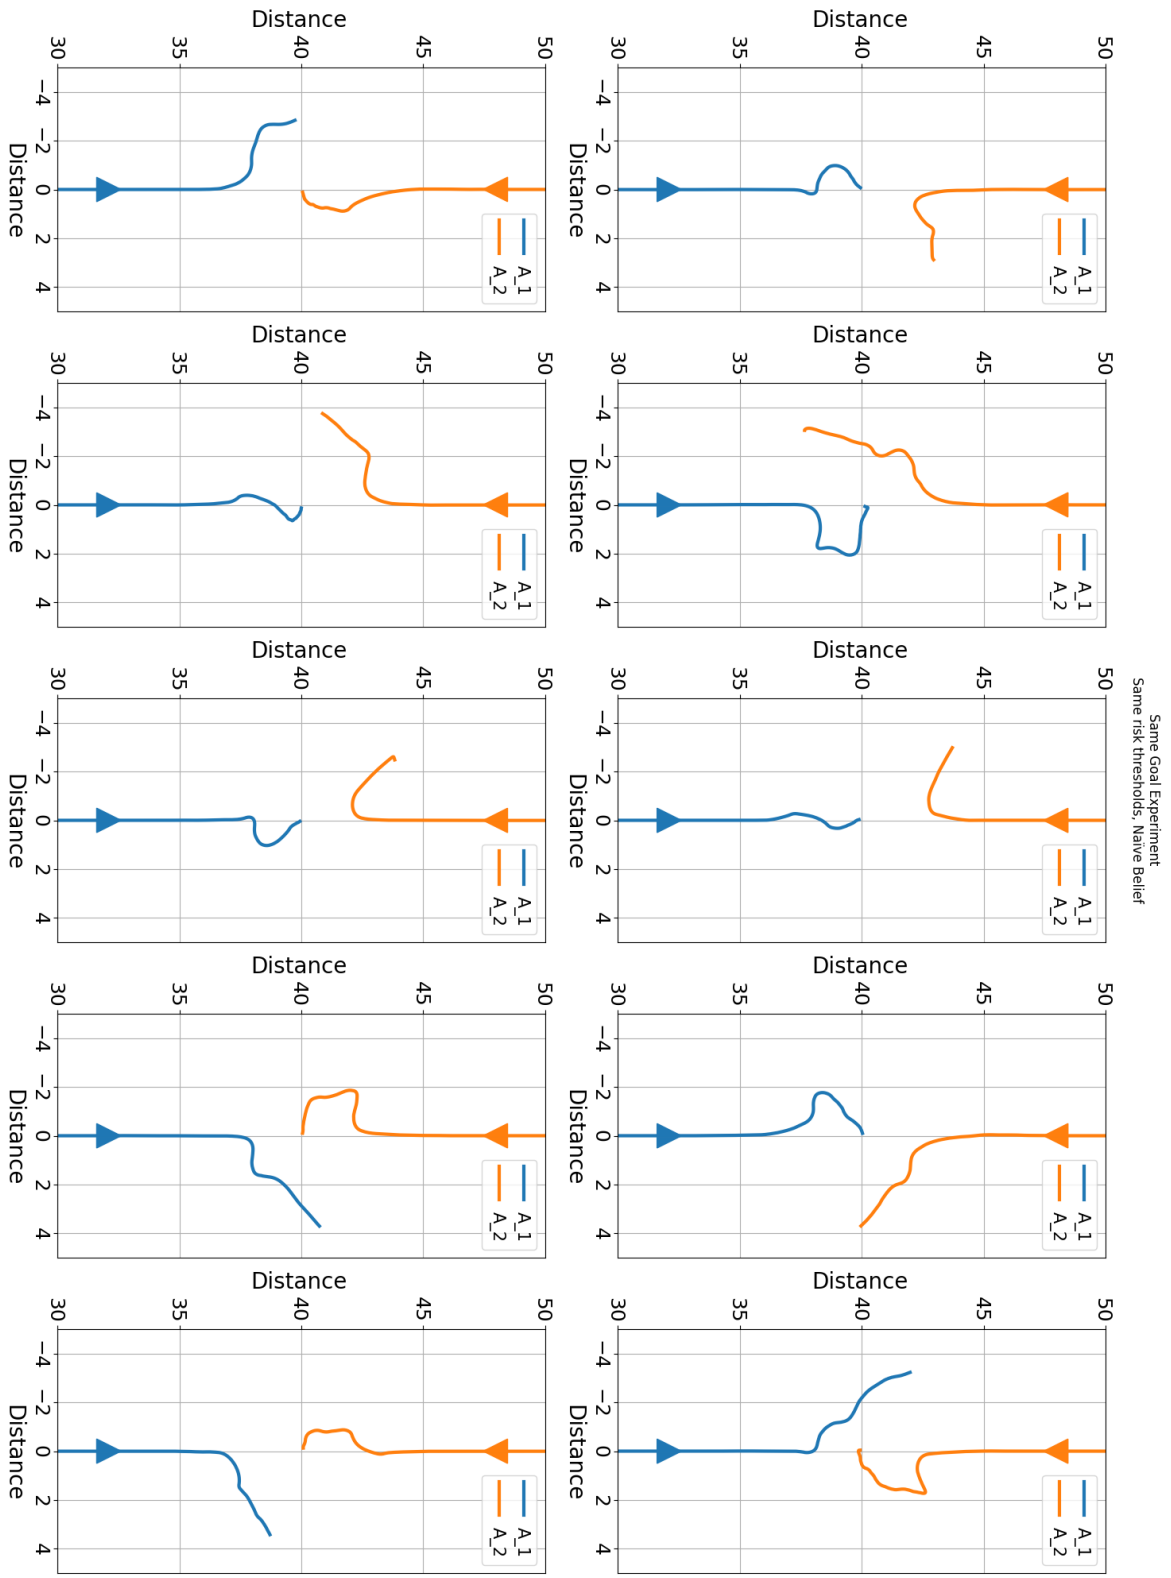


Figure B.28: Experiment: Same Goal, Risk threshold: Equal, Belief: Naïve, magnified

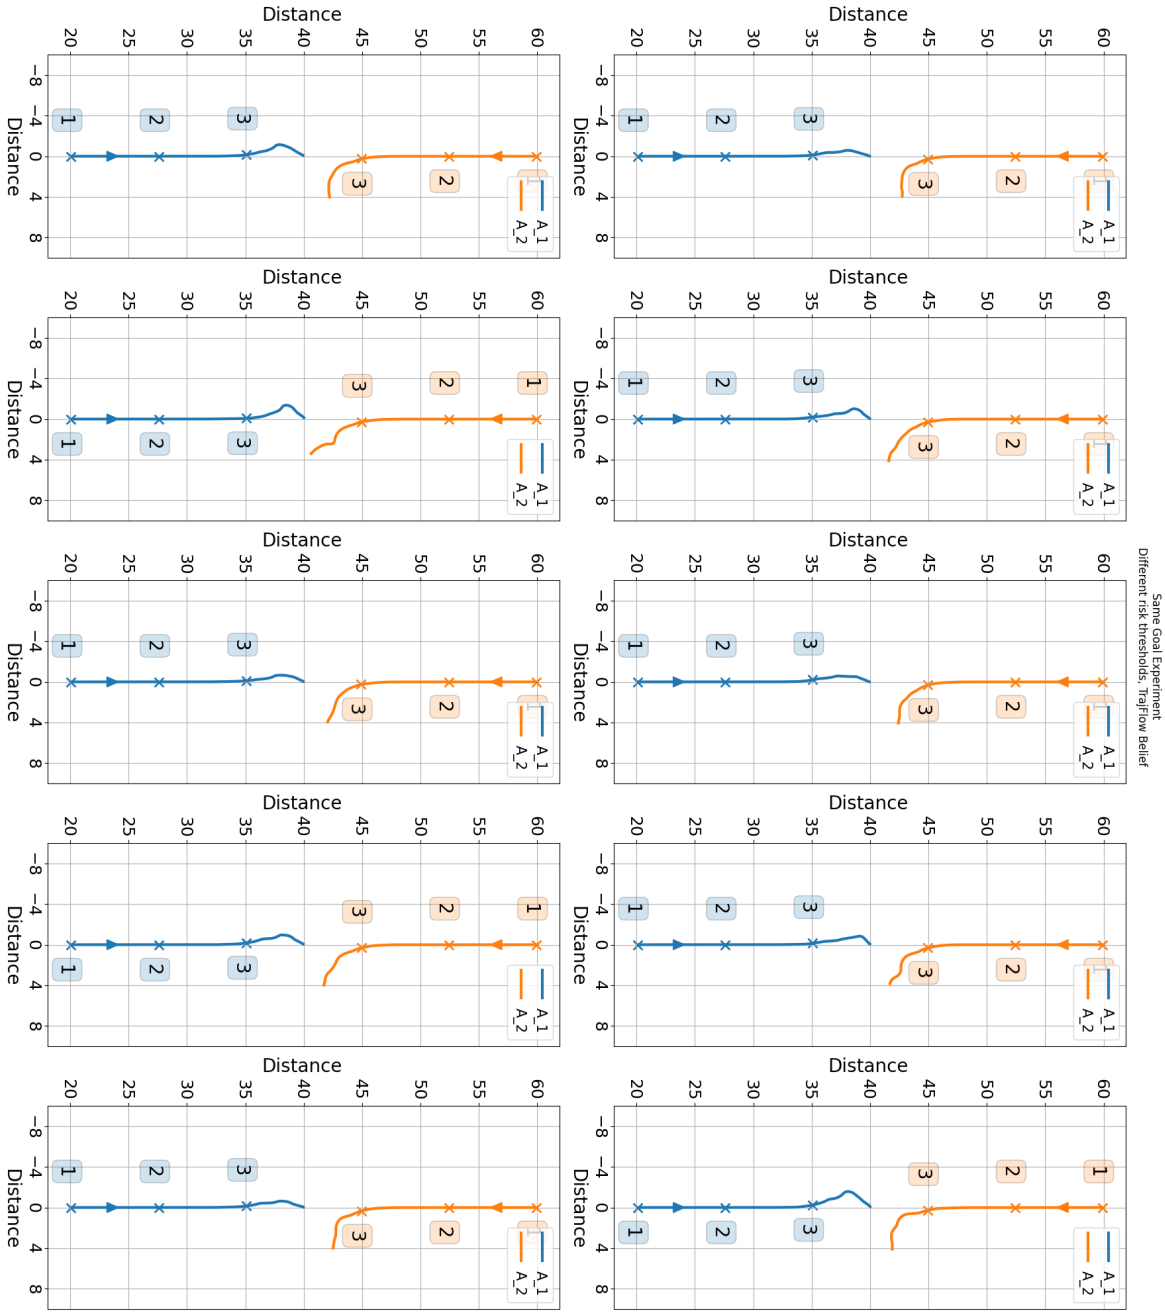


Figure B.29: Experiment: Same Goal, Risk threshold: Different, Belief: TrajFlow

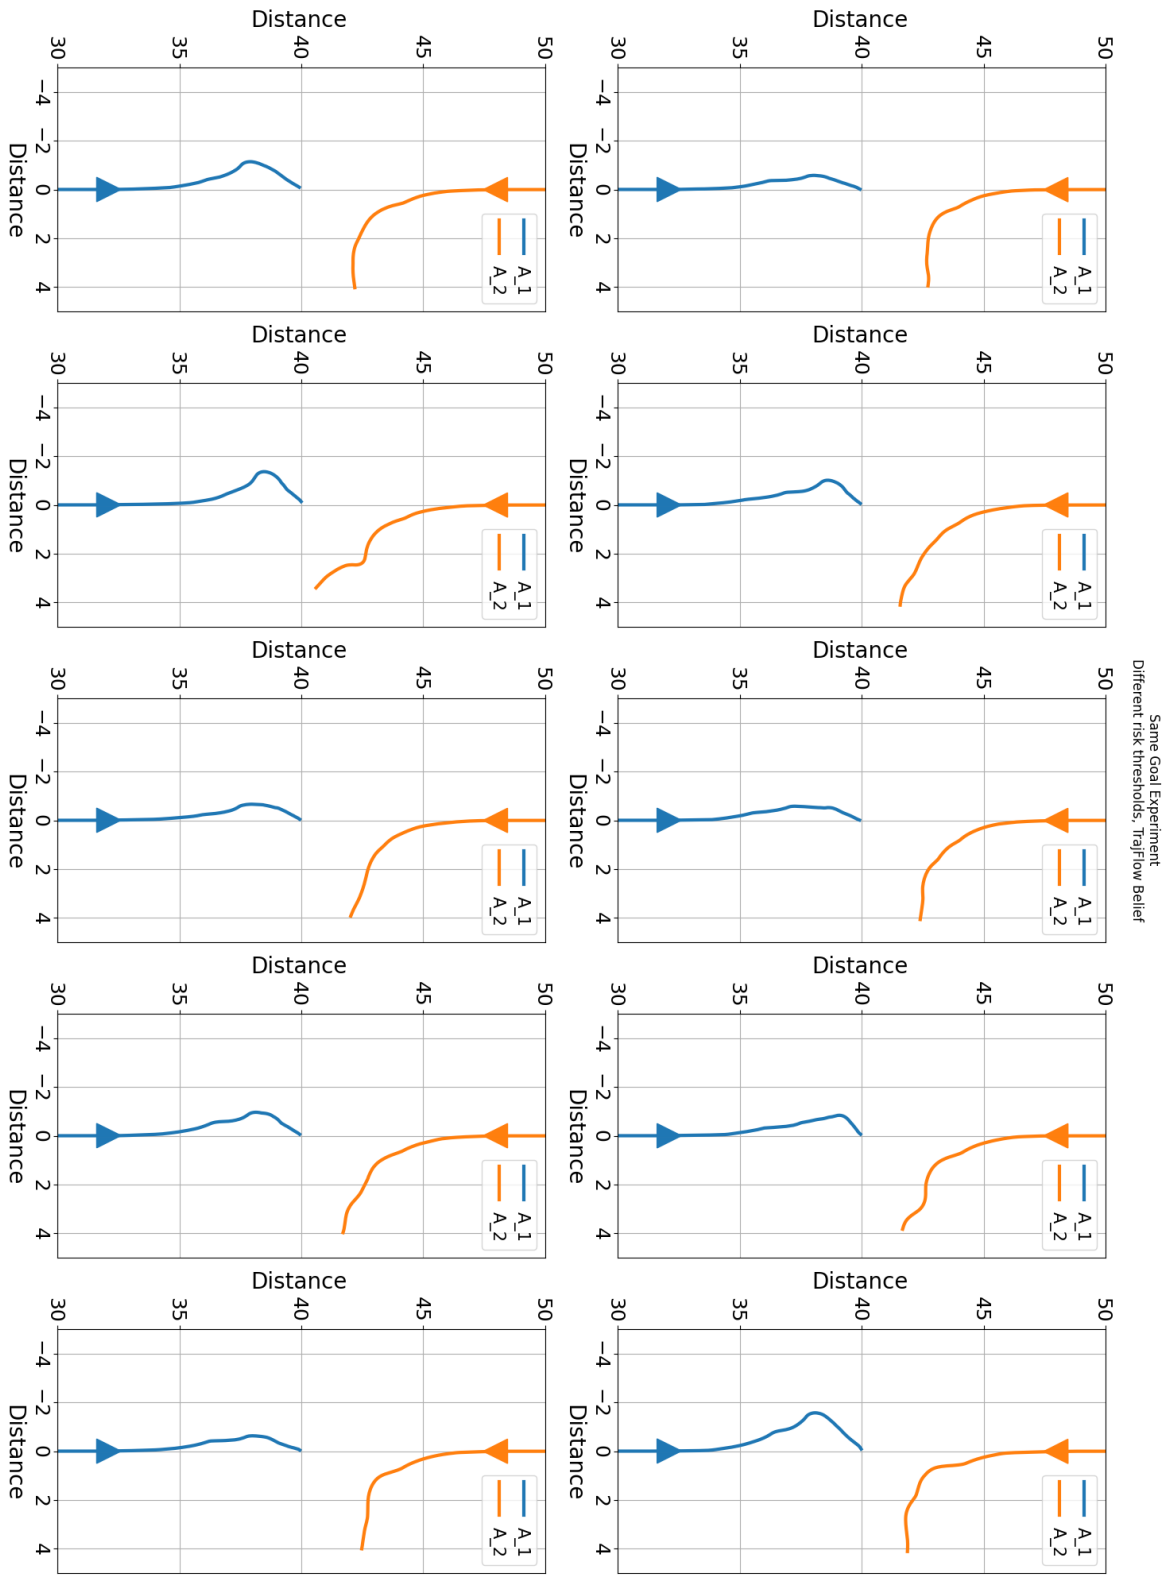


Figure B.30: Experiment: Same Goal, Risk threshold: Different, Belief: TrajFlow, magnified

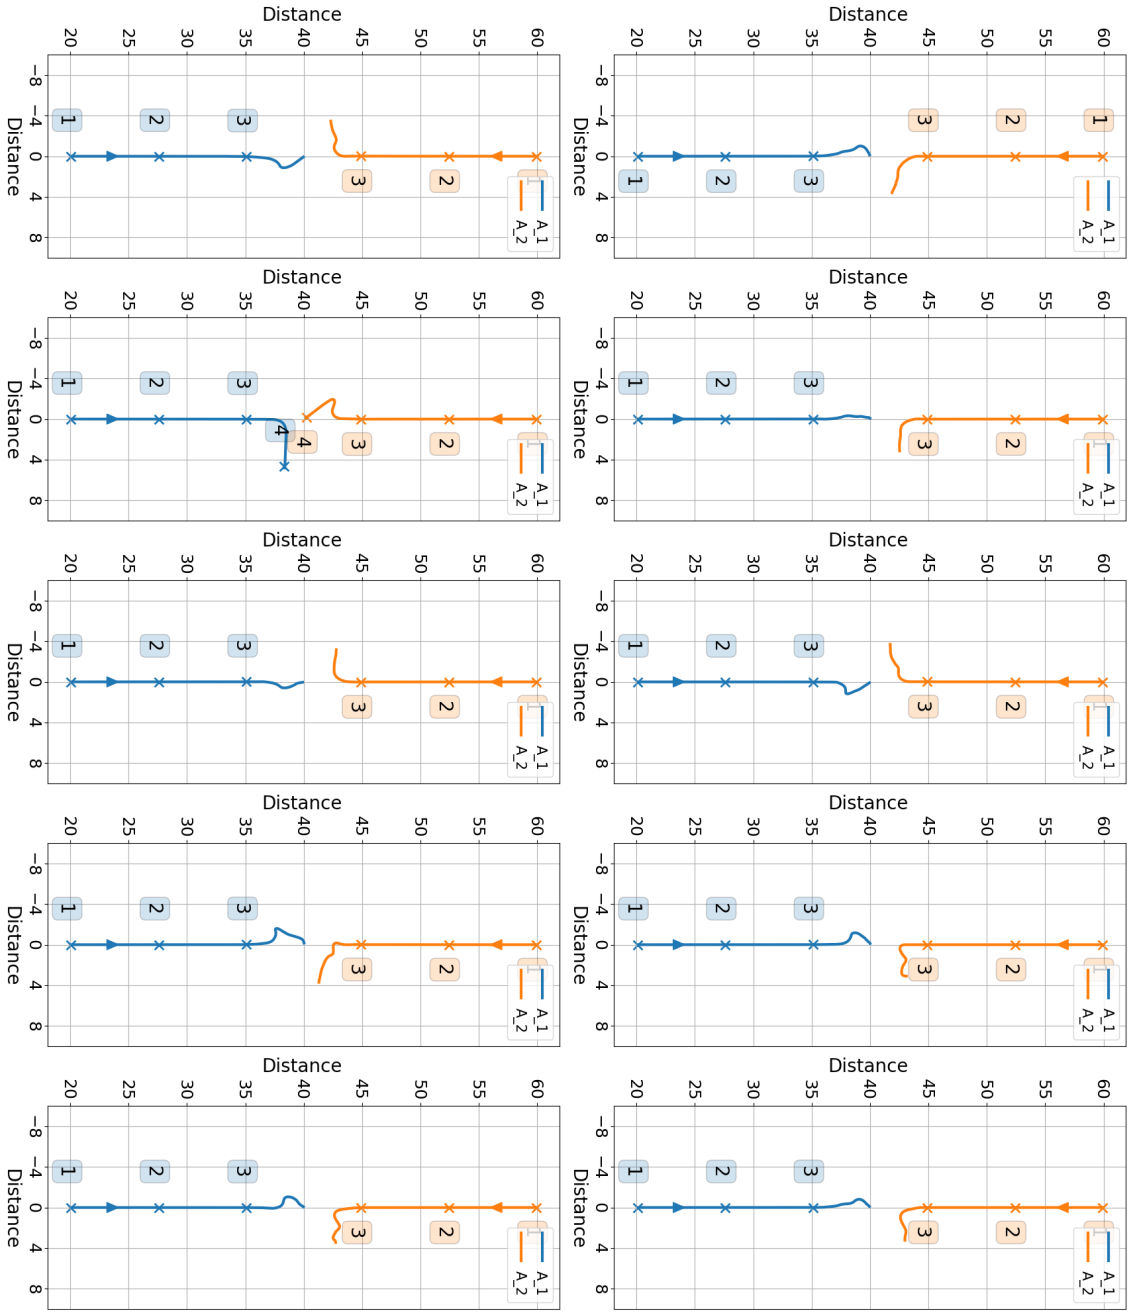


Figure B.31: Experiment: Same Goal, Risk threshold: Different, Belief: Naïve

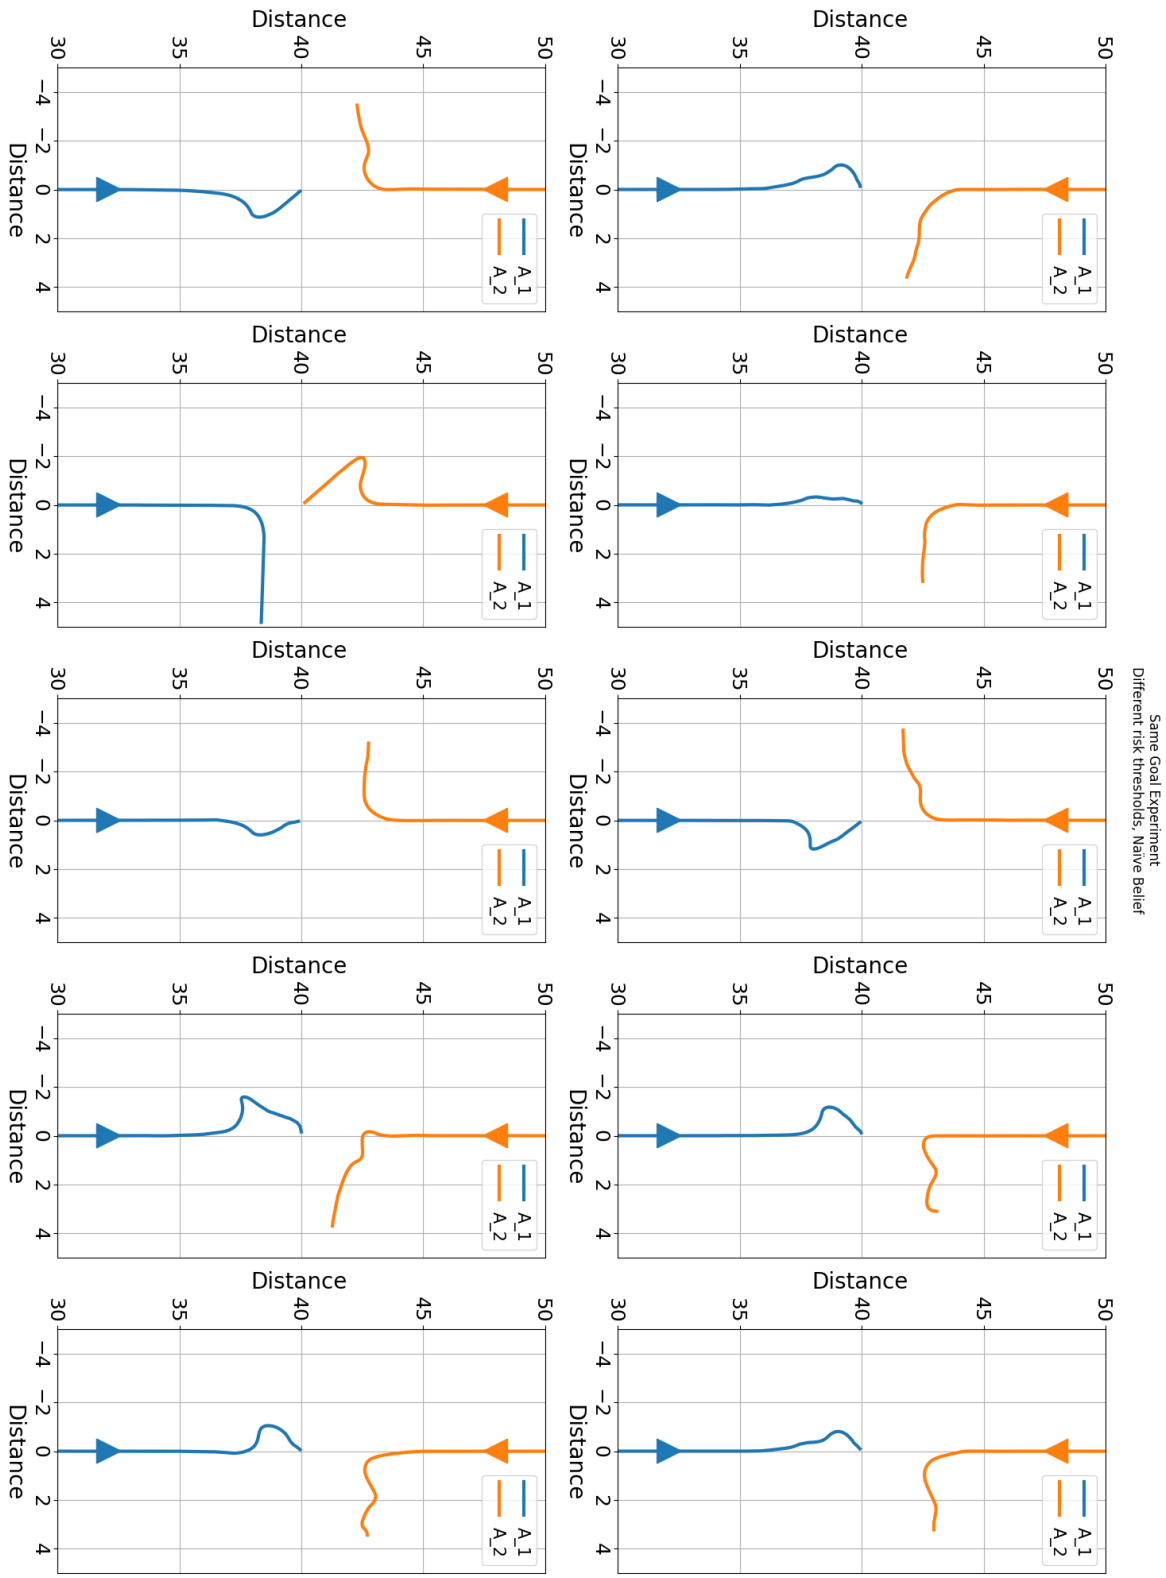


Figure B.32: Experiment: Same Goal, Risk threshold: Different, Belief: Naïve, magnified

# C

# Quantitative Results Experiment

# C.1. Metric: Maximum deviation

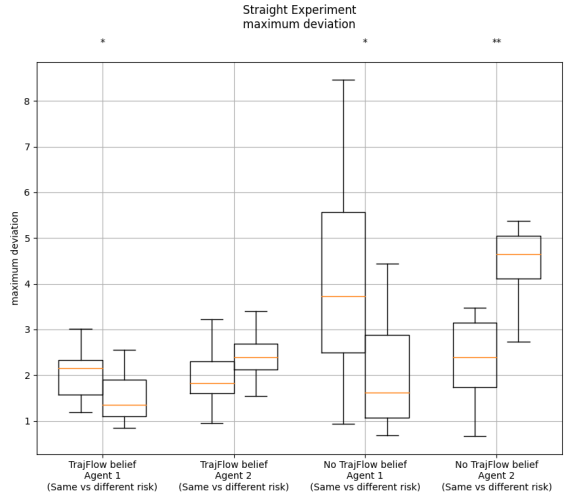


(a) Simulation: straight

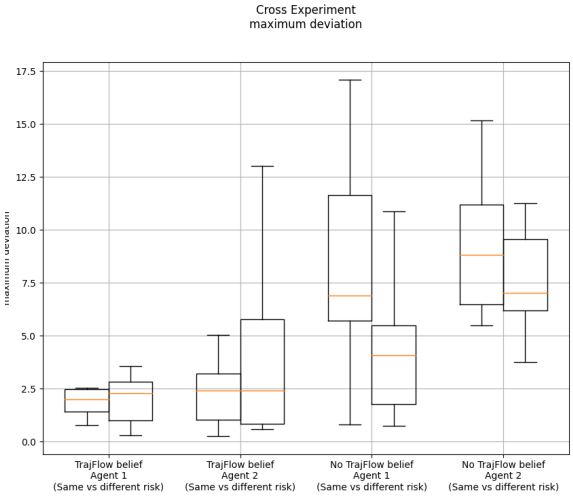


(b) Simulation: cross

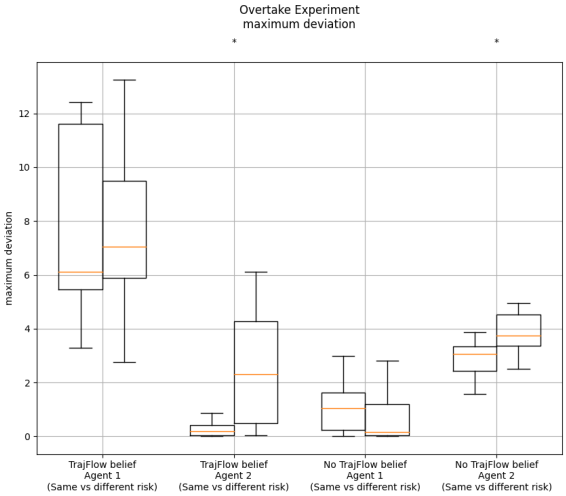


(c) Simulation: overtake

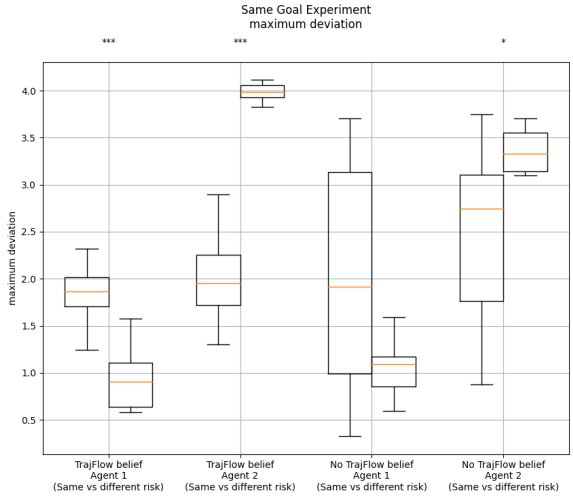


(d) Simulation: same goal

Figure C.1: Metric: maximum deviation

# C.2. Metric: Moment of deviation

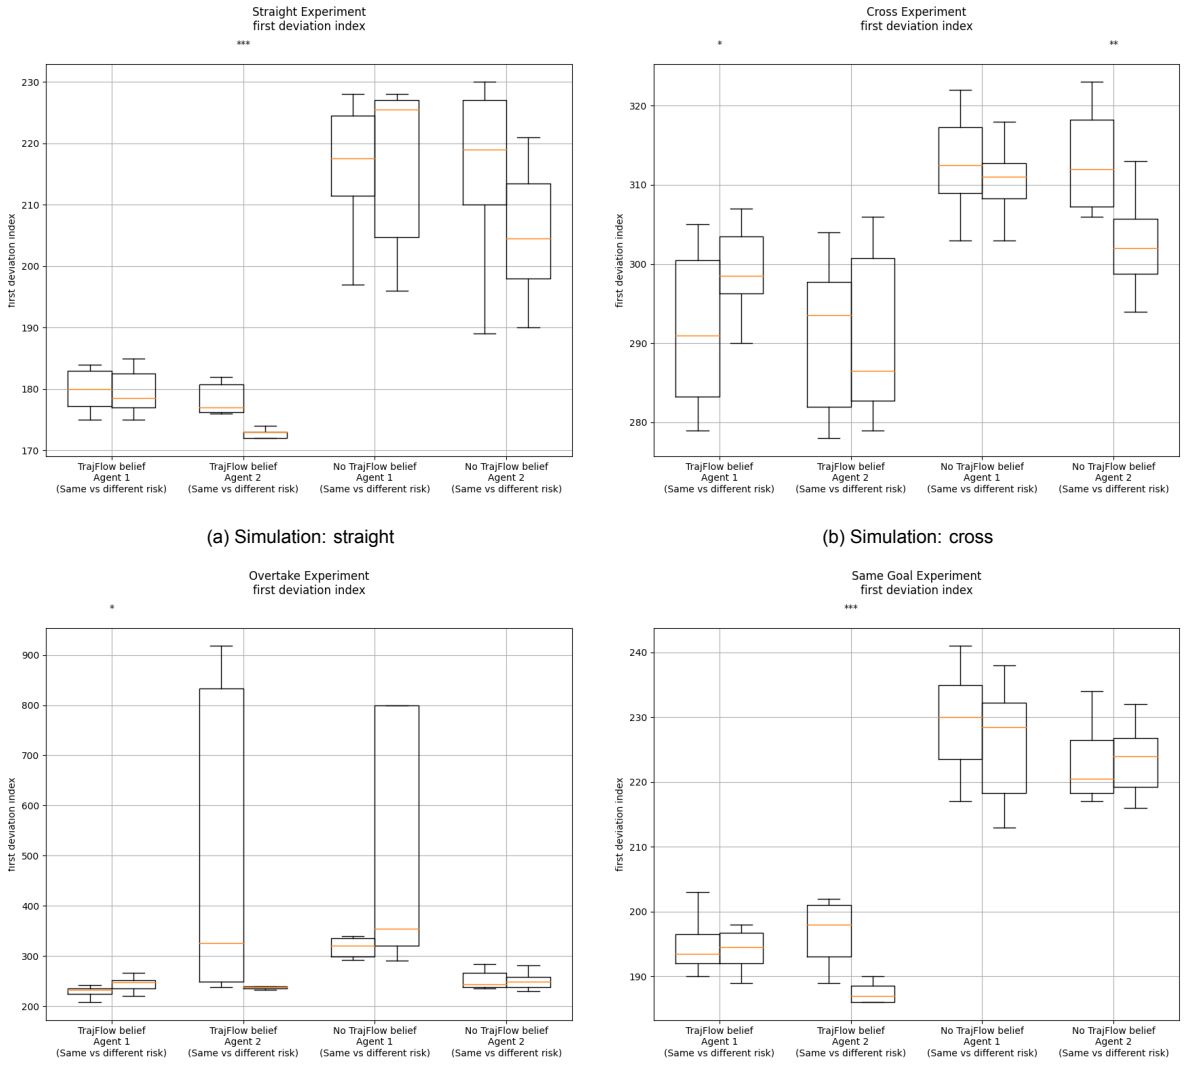


(a) Simulation: straight

(b) Simulation: cross

(c) Simulation: overtake

(d) Simulation: same goal

Figure C.2: Metric: moment of deviation