Bayesian Optimization of Road Pricing using Agent-based Mobility Simulation

Dimitrios Argyros*[1], Renming Liu[1], Ravi Seshadri[1], Filipe Rodrigues[1], and Carlos Lima Azevedo[1]

[1]Department of Technology, Management and Economics, Technical University of Denmark, Denmark

Short summary

Road pricing policies are frequently debated but not widely adopted. Tools for designing near practice-ready policies are still missing, especially considering the complex dynamics between the different levels of traveller decision-making and the networks' performance. We couple an agent- and activity-driven mobility simulator with a Bayesian Optimization (BO) framework for designing optimal road pricing policy in a daily mobility and transportation network system. We extend the literature with a BO-framework application to distance-based road pricing under a departure-time and route-choice sensitive demand model combined with a detailed mesoscopic network. We then tested a general BO and a recently proposed contextual BO algorithm for SimMobility and computational performance. Both identified a similar optimum distance-based pricing, with the second being more computationally efficient. Nonetheless, iterations number, increasing search space and dimensionality could limit their performance. Lastly, the effects of the identified policy were analyzed by leveraging the outcome capabilities of SimMobility.

Keywords: agent-based modelling, Bayesian optimization, congestion pricing, machine-learning, SimMobility

1 Introduction

Numerous strategies have been explored to address traffic congestion and its associated repercussions, which include increased air pollution, accidents, and negative impacts on the overall city's quality of life. Congestion pricing control constitutes one of the main traffic management measures that have been researched but not widely applied during the past decades (Wang et al., 2022; de Palma & Lindsey, 2011). However, certain implementations have produced positive results, such as the Stockholm (Börjesson et al., 2012), and the Milan case (Gibson & Carnovale, 2015). As initially stated by Pigou (1920), congestion pricing serves as a technique for internalising the cost of externalities caused by road users. Since then, various tolling schemes, have been researched and implemented over the years. Regarding the classification of tolls, this occurs according to numerous factors. These can be the kind of scheme (e.g., facility-based, area-based, or distance-based), the level to which charges fluctuate across periods, additional characteristics of toll variation, and equipment (de Palma & Lindsey, 2011).

The most typical price optimisation approaches are either with nonlinear programming or a bi-level optimisation setup, both of which entail a high degree of complexity (NP-hard). Moreover, because of the practical implications they face, the first-best solution is generally used as a benchmark in many studies (Verhoef, 2005; de Palma et al., 2005). As a result, second-best prices, have been most researched, with commonly used methods for calculating prices being heuristics, meta-heuristics, approximations, and trial and error (Verhoef, 2005; Ekstrtopm et al., 2009; Luo et al., 2019).

Bayesian optimisation (BO) is a comparatively recent pricing optimisation approach. Since the pricing problem is NP-hard and does not have simple analytical solutions, black-box optimisation approaches are often used, and the BO has been demonstrated to be a promising one. First, Zhong et al. (2021) presented a BO technique for solving a bi-level optimisation issue where the effects of road pricing on both flow times and land usage were investigated. Particularly, an integrated transport assignment and land usage model was employed in the study and was combined with an active learning algorithm that featured the multivariate-multi-objective BO method. In the same vein, Liu et al. (2021) investigated a BO approach combining a trip-based Macroscopic

Fundamental Diagram (MFD) model with distance-based area-based pricing schemes. The toll profile was created based on time-of-day pricing, with the dimensionality impact on the BO under question. As a result, it was numerically discovered that as the search space expands, performance decreases progressively.

Considering the complexity of the problem, and the lack of tools for designing near practise-ready policies, this study aims to develop approaches for optimal design and evaluation of congestion pricing policies. Following the available research and leveraging the current knowledge, the study brings emerging dynamic congestion pricing control formulations closer to practice, by developing and extending state-of-the-art simulation frameworks. Two frameworks were inspired by the working paper of Liu et al. (n.d.), where BO algorithms were tested on a MFD model for the morning commute problem. However, this study uses a different simulation model which provokes higher complexity and differentiation in concepts but provides the possibility for thorough evaluation and analysis of policies. The study also extends the optimal toll query to a full-day mobility problem with multiple trip purposes.

2 Methodology

We assume the decision maker needs to define a distance-based road pricing scheme for an entire network, which can be time-varying. Since we aim at dining the best toll for an average day, the network equilibrium process and the optimization process may interact with each other. In a simpler and theoretically sound mobility model, it was recently shown that adding context data to BO related to the equilibrium process can help in reaching optimal solutions faster (Liu et al. (n.d.)). We test again such a hypothesis in a more complex and closer to reality simulator. In the following paragraphs, we will describe (1) the overall simulation platform, SimMobility; (2) its learning (equilibrium) process; (3) the general BO approach; (4) the contextual BO; and (5) our specific combined framework for both the general and the contextual BO cases.

SimMobility [1], a detailed agent-based mobility simulator, was used in the context of this study. Specifically, SimMobility constitutes a large-scale simulator that includes many mobility-sensitive behavioural models in a multi-scale dimension. It was built from three distinct sub-models: long-term (LT), mid-term (MT), and short-term (ST). The MT model is made up of three interconnected simulators: 1) the pre-day simulator, which calculates the individual daily activity schedule (DAS), 2) the within-day simulator, which simulates departure times and route choice behaviour including en-route behaviour, and 3) the supply mesoscopic simulator, which handles network characteristics and the supply for different modes of transport. The reader is referred to (Adnan et al., 2016; Lu et al., 2015) for more details on SimMobility's models and implementation.

The MT provides two learning approaches, with the one used known as within-day learning. In more detail, an activity schedule is initially given and the default travel time is derived for each link. Then, and for each time a within-day simulation is conducted, the link travel times experienced by each vehicle update the database accordingly. The second, namely day-to-day learning, occurs when the agents' transport mode and their relevant travel information are updated. By iterating the pre-day, a new DAS is produced, and used as input to the within-day one, which consequently allows a learning process (Lu et al., 2015).

As for the BO, its main components are a model of the objective function and an acquisition function. Specifically, it develops a surrogate for the objective and quantifies the uncertainty in that by utilising a Gaussian process (GP) regression, fitting the data, and then selecting the next sampling point based on the acquisition function (Frazier, 2018). GP regression is described by a prior mean function $\mu_0(x)$ and a covariance function $\kappa(x, x')$, where x denotes the input variables. For the simplification of the GP training, $\mu_0(\cdot) = 0$ was used (Rasmussen & Williams, 2005). As for the covariance function, the Marten Kernel, defined below, can be used, where $\Gamma(\cdot)$ is the Gamma function and H_ν is the modified Bessel function. Here $\nu = 5/2$ is used.

$$W_{(x)} \sim GP(\mu_0(x), \kappa(x, x')) \tag{1}$$

$$W_{(x_{n+1})} | W_{(x_n)} \sim \mathcal{N}(\mu(x_{n+1}), \sigma^2(x_{n+1})) \tag{2}$$

[1]https://github.com/smart-fm/simmobility-prod

$$k(x,x') = \frac{2^{\nu-1}}{\Gamma(\nu)}\left(\frac{\sqrt{2\nu}}{l}||x-x'||\right)^{\nu} H_{\nu}\left(\frac{\sqrt{2\nu}}{l}||x-x'||\right) \tag{3}$$

The acquisition function is used to determine the next point to assess, given the posterior mean function and the variance of the GP. The upper confidence bound (UCB) is a widely used one (Srinivas et al., 2012) where ρ is a hyperparameter that determines the ratio of exploration to exploitation, with a bigger value indicating higher exploration. Here $\rho = 2$ was chosen.

$$n_{(UCB)}(x;\rho) = -\mu(x) + \rho\sigma(x) \tag{4}$$

$$x_{i+1} = \underset{x}{\operatorname{argmax}}\, n_{(UCB)}(x) \tag{5}$$

Regarding the contextual BO, Krause & Ong (2011) presented that the BO can take into consideration different environments through the usage of contextual variables. Thus, a contextual Gaussian process (CGP) was proposed. Here, the contextual variable which is used for the CGP is the day d. In that setting, and as the product of two Merten kernels is still a Merten kernel, the composite kernel occurs accordingly.

$$\overline{k}(x,d),(x',d') = k(x,x') * k(d,d') \tag{6}$$

As for the evaluation of the system, the main key performance indicators used were the travel time index (TTI), the Social welfare (SW) and the consumer surplus (CS). For the computation, only the trips by car are considered. TTI is the ratio of average travel time to free flow travel time. It is a metric of congestion that can be used to assess overall network performance and, here, is calculated at the 5-minute level, where $TTobserved_{n_i}$ denotes the travel time observed, $TTfreeflow_{n_i}$ the travel time in free flow settings and $TtripLength_{n_i}$ the total trip length, for the traveller n in the time interval i. SW is used for assessing the overall performance of the system and is computed based on the total of observed individual travel disutilities, where D_n denotes the disutility, $TravelCost_n$ the travel cost and β_{Cost_n} the cost coefficient for the traveller n. The CS of the system is calculated below and represents the total travellers' expenses.

$$TTI_i = \frac{\sum_{n_i=1}^{N_i} \frac{(TTobserved_{n_i} * TtripLength_{n_i})}{TTfreeflow_{n_i}}}{\sum_{n_i=1}^{N_i} TtripLength_{n_i}} \tag{7}$$

$$SW = \sum_{n=1}^{N} \frac{(D_n - TravelCost_n * \beta_{Cost_n})}{\beta_{Cost_n}} \tag{8}$$

$$CS = \sum_{n=1}^{N} \frac{D_n}{\beta_{Cost_n}} \tag{9}$$

Regarding the frameworks, both extend the SimMobility simulator and interact with it by updating its database; thus, their integration happens. Their objective function is the SW, whereas the acquisition function predicts the (distance-based) toll's control variables. Noteworthy is that the tolling rates are imputed to the database at 5-minute intervals based on a multi-modal Gaussian distribution resulting from the BO algorithms, where the number of modes corresponds to the number of network traffic peaks (p). In this setting, the amplitude (A_p), represents the height/highest price, the mean (μ_p) indicates the time of the highest price, and the standard deviation (σ_p) defines the width of the profile.

As for framework 1 (fig. 1), the simulation starts with a sampled toll and when the system has reached equilibrium, the BO forecasts the next test points/toll. Then the database is updated, by initialising the link travel times at no-toll conditions and incorporating the new toll prices, and the within-day simulator is iterated until the equilibrium is reached. The process is repeated until n the GP has converged by reaching the optimum toll profile.

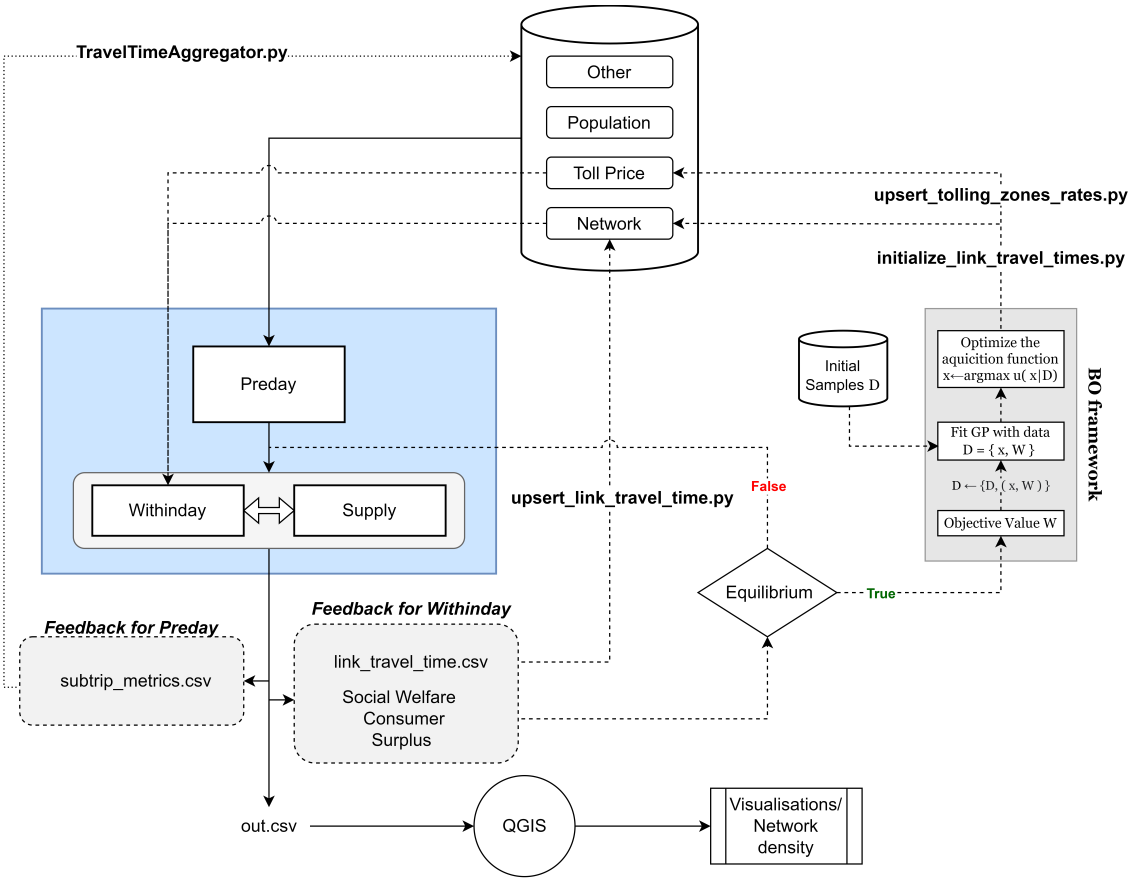

Figure 1: Framework 1. Bayesian optimisation and SimMobility

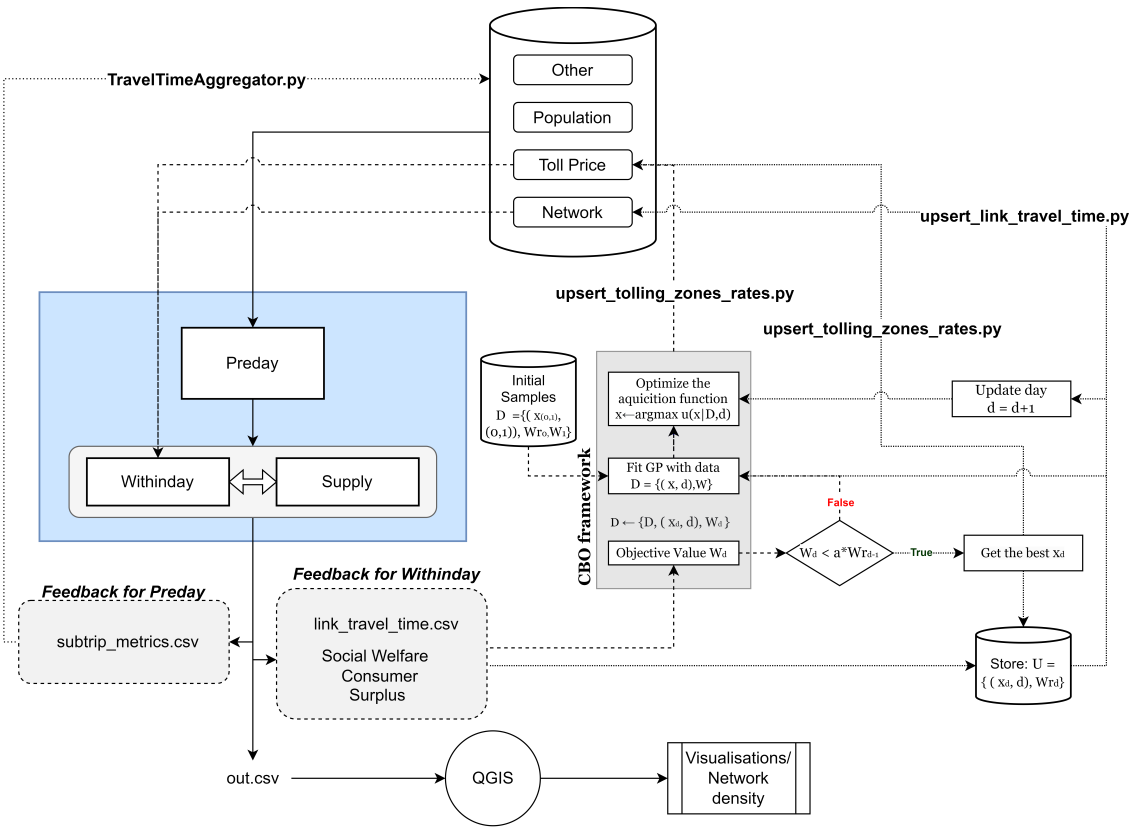

Figure 2: Framework 2. Contextual Bayesian optimisation and SimMobility

Framework 2 (fig. 2) follows a similar design but aims to offer greater flexibility as it can alter the day, based on the condition that the day's SW is less than the preceding one multiplied by a scaling factor (a). In more detail, the framework considers as initial inputs the SW of day zero (SW of the no-toll case), and a few input points for day 1. Then the contextual BO proposes the next toll for the same day and SimMobility calculates its SW. The same procedure continues until a toll which gives a better SW than the previous day's has been identified. Then the realised SW is calculated for the day and stored. Simultaneously, the framework updates the perceived link travel times of the network. Then, it updates the day and the algorithm proposes the first toll for the new day. This approach would be repeated until an overall equilibrium is attained for both the system and the GP.

Lastly, an initial set of points is essential for the initialization of the BO algorithms. Thus, a space-filling experiment design which uses Latin hypercube sampling (LHS) was included in both frameworks (Mckay et al., 2000).

3 Results and discussion

SimMobility's built-in prototypical city was used for the experiments. Its network has 254 links, which correspond to 286 segments. The total length of the links is 279 km with an average maximum speed of 65.27 km/h and a standard deviation of 12.77 respectively. The number of nodes of the network is 95 with 1918 turning intersections. Lastly, the city can be separated into 24 zones. Also, a fixed demand of 19000 travellers, corresponding to 51071 trips, were employed. Moreover, to attain a congested condition of the system, the network's initial capacity was reduced to 10%, despite a few bottlenecks, which remain at 30% of the original capacity. Based on the new capacity, the average critical density of the system per link is at 6,62 passenger-car units(PCU)/km with a standard deviation calculated at 2,03.

First, an experimental design took place to verify and evaluate the frameworks. The frameworks set to design a distance-based toll policy with an upper payment bound (UB) at $10. The purpose of this experiment was to prove that the developed frameworks are adequate to design and assess a toll which optimises the system's SW. On top of that, the comparison of the frameworks' performance was set under question. Moreover, it should be noted that for the experiments, the pre-day simulation is inactive and considered that it has already provided the modal split and the respective DAS file.

Table 1: Optimisation bounds

Bounds	μ_1 (5-min)	σ_1 (5-min)	A_1 ($/m)	μ_2 (5-min)	σ_2 (5-min)	A_2 ($/m)
Upper	108	24	0.003	234	24	0.003
Lower	84	3	0.00001	210	3	0.00001

It has been found that both frameworks can indeed design a toll that raises the SW of the systems. Also, the fact that both frameworks reached a similar toll profile is a clear sign that the global optimum, within the respective constraints, has been discovered. Regarding their performance (see fig. 3 and fig. 4), the number of iterations required for framework 1 was for the 12 sampled points 240 ($\sim$70 hours), and for the algorithm, until the optimum been found after 15 repetitions, was 300 ($\sim$88 hours). However, considering the number of needed iterations for the system to reach stationary, 126 and 150 iterations, respectively, are determined as a "fair" number. As for CBO, the setting immediately offers many advantages as just 5 iterations were required for initialization. The framework then ran with varying iterations per day, reaching the ideal toll region after $\sim$110 iterations ($\sim$32 hours).

As for the designed toll, the one with the biggest SW provided by framework 2 (μ_1: 108.0, σ_1: 14.39, A_1: 0.003, μ_2: 234.00, σ_2: 24.00, A_2: 0.00001) was selected for further analysis. The toll has a very high peak, similar to the morning traffic peak, which has pushed some trips earlier in the morning. In this way, the vehicle accumulation and the TTI have decreased compared to the no-toll case, from $\sim$1200 to $\sim$900 and from 1,65 to 1,35 respectively (see fig. 5). As for the afternoon peak, CBO has not defined a toll profile that affects the peak's characteristics, while the amplitude was set at the lower bound. Regarding SW, this was calculated to $92611 implying a 5,5% improvement

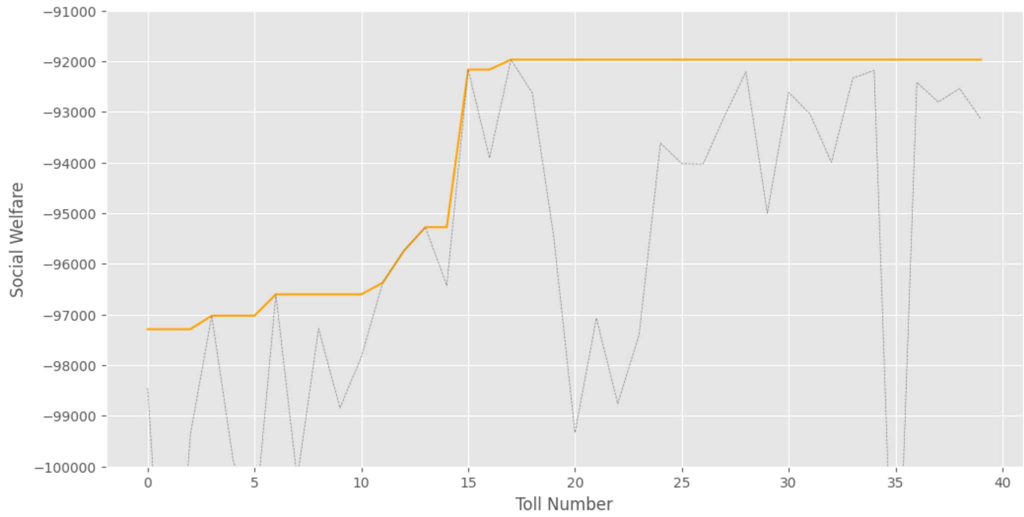

(a) Framework 1

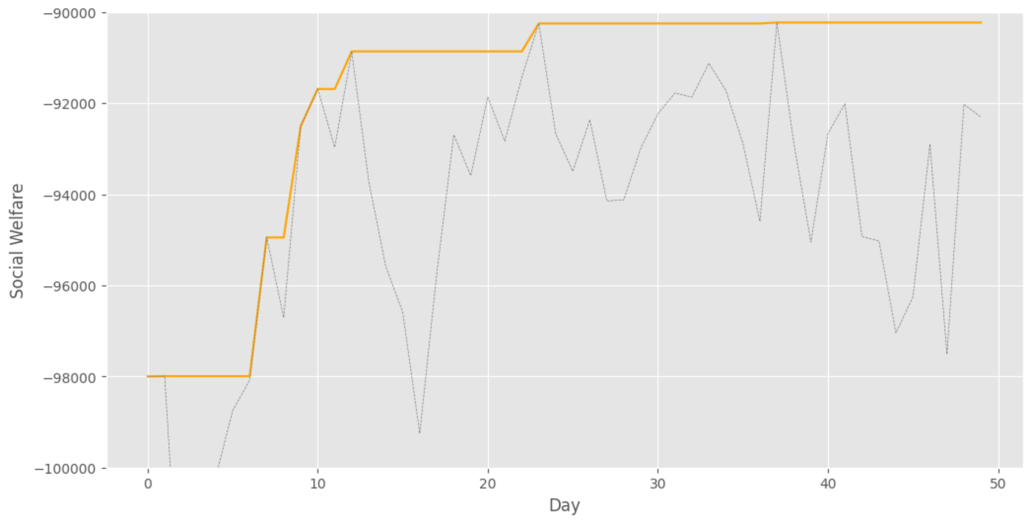

(b) Framework 2

Figure 3: Best social welfare ($) (yellow line) and social welfare ($) evolution for tested tolls (grey line)

Table 2: Base case and optimum toll case metrics

	Base case		**Toll case**		**MD**
	Mean	SD	Mean	SD	%
Social welfare ($)	-98006.82	2707.91	-92611.51	927.19	5.51
Consumer surplus ($)	-98006.82	2707.91	-231406.13	963.59	-136.11
Trip avg. travel time (sec)	458.38	6.23	410.07	2.43	10.54
Trip avg. scheduled delay (sec)	332.74	5.44	515.95	1.62	-55.06
Trip avg. paid toll ($)	0.00	0.00	3.31	0.00	-

compared to the base case (see table 2). Similarly, the mean travel time improvement is around 10,5%, where the average scheduled delay has become 55% longer. Moreover, the most gains in travel time have been identified in the southeast part of the city, while the bigger CS losses were presented in the northern region. Furthermore, the link densities during the morning peak hour (08:00–09:00 AM) have been decreased. Specifically, the number of links during the peak hour that presented a density bigger than the average critical density of the system, decreased from 64 to 39 due to the implementation of the toll. Similarly, the average system density in this period declined from 2.3 to 1.9 PCU/km/lane. Additionally, for travellers, it has been calculated that in the equilibrium state, ∼20 % pays the maximum limit of $10 and ∼55% less than $1 (see fig. 6).

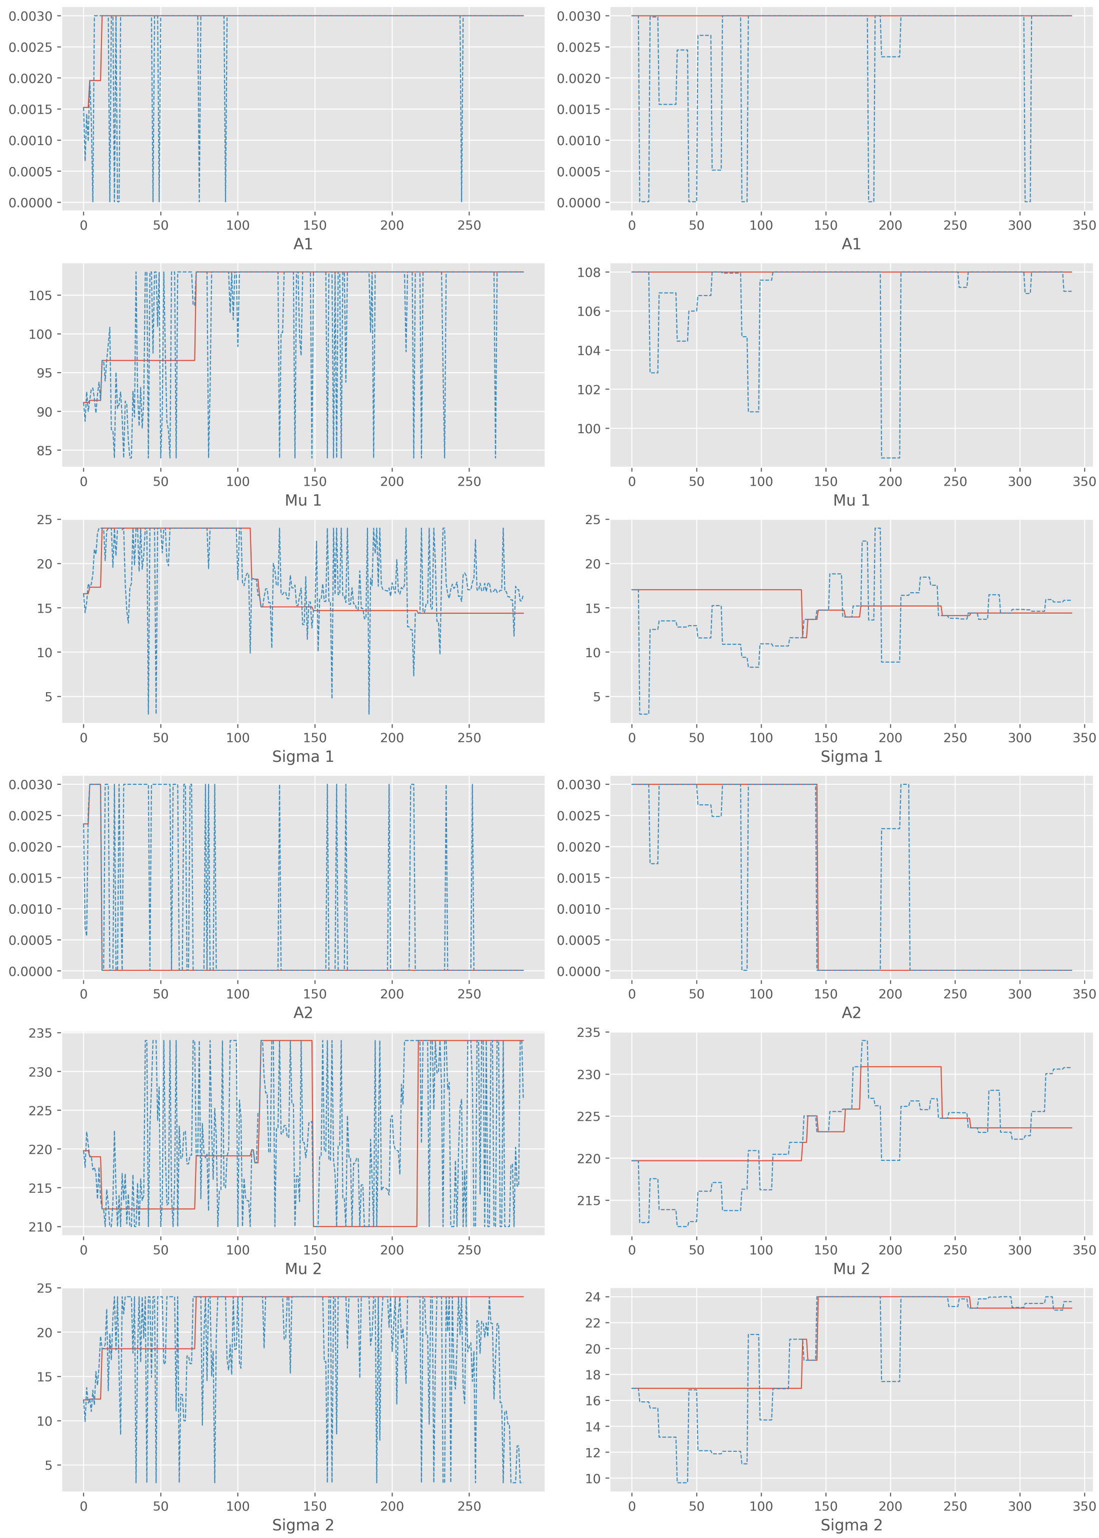

Figure 4: Selected parameters across all iterations for both frameworks.
Note: No sampled data included. For framework 1 only the iterations until stationary are included. The red line indicates the parameter for best social welfare. The first column refers to framework 2 and the second to framework 1

Lastly, in another direction, a shortcoming of the toll was also identified, as travellers departing after the peak hour (10:00-11:00 AM) faced a high toll charge and added scheduled delay due to it. That could be mainly reasoned to the system's demand which is comparatively low during this period and thus the applied toll during this period does not affect significantly the SW. That consequently may prevent the frameworks to design a better solution.

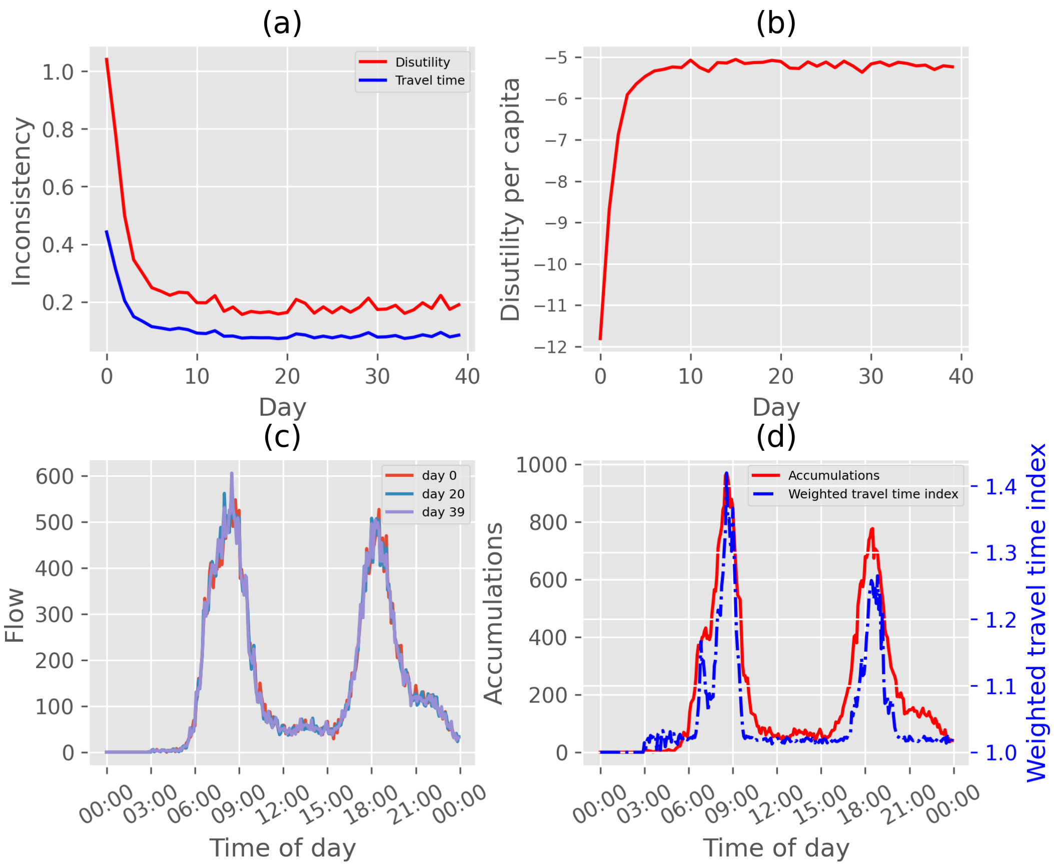

Figure 5: Optimum toll case. Simulation and congestion graphs.
Note: The accumulation and TTI are based on the last iteration

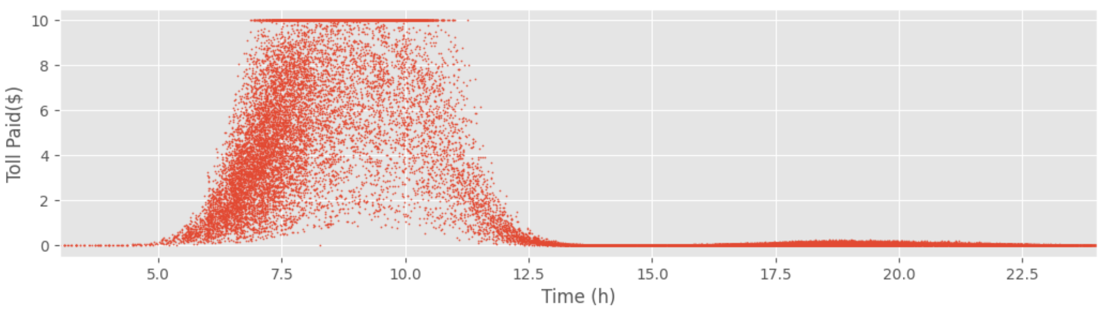

Figure 6: Optimum toll case. Amount of toll paid ($) per user in equilibrium state
Note: The graph is based on the last iteration

In other experiments, framework 2 attempted to determine whether there was a genuine distance-based toll. It was found that the amplitudes converged to zero since the framework could not find a toll that enhances the SW. Nonetheless, when the framework was asked to find a very low price that improves the SW, it was not confident of any solution after a sufficient number of iterations. This limitation may arise from the fact that the differences in the SW system shown for such a small toll could not be significantly different from the no-toll case. Also, considering the quite big system variance, with a standard deviation at 3%, this might be misleading. Similarly, framework 2 was set to find a better solution for the SW, while a larger search space in terms of boundaries and dimensions was set. Nonetheless, that was not possible in 30 days. The fact that in these 30 days, it ran only for ~ 90 iterations in total may be one of the causes of this failure. In addition to the number of iterations, the acquisition weight could also have an impact on the solution as the search space was expanded.

4 Conclusions

During this research, two frameworks combining a detailed agent-based simulator, namely SimMobility, and BO algorithms, were developed and tested. Their main goal was to determine the optimum road pricing policy of a day by optimising the system's SW. It was found that both frameworks were able to pinpoint the optimum distance-based toll design with UB at $10. Nevertheless, in terms of performance, the second framework proved that needs much fewer iterations.

The advantages of the analysis of a road pricing policy through a detailed agent-based simulation were also showcased. First, the policy was able to improve the SW of the system as well as the TTI during the peak period. Also, by diving into more disaggregated metrics, the identification of improvements in link-level density, during the peak period, was possible. In the same vein, by using a zonal level analysis, the regions with bigger travel time gains and CS losses per traveller, which were the city's southeast and northern parts respectively, were identified.

Moreover, framework 2 found that there is no genuine distance-based toll that improves the SW in this environment setup. However, the framework's inability to identify a solution, when the amplitude's search space significantly decreased, was identified. This may be due to the number of iterations combined with the system's inherent variability. Finally, framework 2 could not locate the same solution or discover a better toll profile when faced with an expansion of the size and dimensions of the search space. Thus, more iterations or a higher acquisition weight may be necessary to resolve this issue.

References

Adnan, M., Pereira, F., Azevedo, C. L., Basak, K., Lovric, M., Raveau, S., ... Ben-Akiva, M. (2016). Simmobility: A multi-scale integrated agent-based simulation platform..

Börjesson, M., Eliasson, J., Hugosson, M. B., & Brundell-Freij, K. (2012, 3). The stockholm congestion charges-5 years on. effects, acceptability and lessons learnt. *Transport Policy*, *20*, 1-12. doi: 10.1016/j.tranpol.2011.11.001

de Palma, A., Kilani, M., & Lindsey, R. (2005, 8). Congestion pricing on a road network: A study using the dynamic equilibrium simulator metropolis. *Transportation Research Part A: Policy and Practice*, *39*, 588-611. Retrieved from https://linkinghub.elsevier.com/retrieve/pii/S0965856405000388 doi: 10.1016/j.tra.2005.02.018

de Palma, A., & Lindsey, R. (2011, 12). Traffic congestion pricing methodologies and technologies. *Transportation Research Part C: Emerging Technologies*, *19*, 1377-1399. Retrieved from https://linkinghub.elsevier.com/retrieve/pii/S0968090X11000362 doi: 10.1016/j.trc.2011.02.010

Ekstrtopm, J., Engelson, L., & Rydergren, C. (2009). Heuristic algorithms for a second-best congestion pricing problem. *NETNOMICS: Economic Research and Electronic Networking*, *10*, 85-102. doi: 10.1007/s11066-008-9019-9

Frazier, P. I. (2018, 7). A tutorial on bayesian optimization. Retrieved from http://arxiv.org/abs/1807.02811

Gibson, M., & Carnovale, M. (2015, 9). The effects of road pricing on driver behavior and air pollution. *Journal of Urban Economics*, *89*, 62-73. doi: 10.1016/j.jue.2015.06.005

Krause, A., & Ong, C. S. (2011). *Contextual gaussian process bandit optimization.* Retrieved from https://www.researchgate.net/publication/229310653

Liu, R., Jiang, Y., & Azevedo, C. L. (2021, 6). Bayesian optimization of area-based road pricing. Institute of Electrical and Electronics Engineers Inc. doi: 10.1109/MT-ITS49943.2021.9529326

Liu, R., Jiang, Y., Seshadri, R., Azevedo, C. L., & Ben-Akiva, M. (n.d.). *Contextual bayesian optimization of congestion pricing with day-to-day dynamics.* (unpublished)

Lu, Y., Adnan, M., Basak, K., Pereira, F. C., Ben-Akiva, M., Carrion, C., ... Saber, V. H. (2015). Simmobility mid-term simulator: A state of the art integrated agent based demand and supply model.. Retrieved from https://www.researchgate.net/publication/279223479

Luo, Q. Y., Zhu, J. X., Guan, X. Y., & Jia, H. F. (2019). Optimal cordon-based congestion pricing considering feasible cordons and congestion transfer. *IEEE Access*, *7*, 93545-93557. doi: 10.1109/ACCESS.2019.2928368

Mckay, M. D., Beckman, R. J., & Conover, W. J. (2000, 2). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, *42*, 55-61. Retrieved from `http://www.tandfonline.com/doi/abs/10.1080/00401706.2000.10485979` doi: 10.1080/00401706.2000.10485979

Pigou, A. C. (1920). *The economics of welfare*. Macmillan.

Rasmussen, C. E., & Williams, C. K. I. (2005). *Gaussian processes for machine learning*. The MIT Press. Retrieved from `https://direct.mit.edu/books/book/2320/gaussian-processes-for-machine-learning` doi: 10.7551/mitpress/3206.001.0001

Srinivas, N., Krause, A., Kakade, S. M., & Seeger, M. W. (2012, 5). Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, *58*, 3250-3265. doi: 10.1109/TIT.2011.2182033

Verhoef, E. T. (2005, 11). Second-best congestion pricing schemes in the monocentric city. *Journal of Urban Economics*, *58*, 367-388. Retrieved from `https://linkinghub.elsevier.com/retrieve/pii/S0094119005000495` doi: 10.1016/j.jue.2005.06.003

Wang, Y., Geng, K., May, A. D., & Zhou, H. (2022, 3). The impact of traffic demand management policy mix on commuter travel choices. *Transport Policy*, *117*, 74-87. doi: 10.1016/j.tranpol.2022.01.002

Zhong, S., Gong, Y., Zhou, Z., Cheng, R., & Xiao, F. (2021, 4). Active learning for multi-objective optimal road congestion pricing considering negative land use effect. *Transportation Research Part C: Emerging Technologies*, *125*. doi: 10.1016/j.trc.2021.103002