

Personalized Incentives with Constrained Regulator’s Budget

Lucas Javaudin^a, Andrea Araldo^b and André de Palma^a

^aTHEMA, CY Cergy Paris Université; ^bTelecom SudParis, Institut Polytechnique de Paris

ARTICLE HISTORY

Compiled November 24, 2023

ABSTRACT

We consider a regulator driving individual choices towards increasing social welfare by providing personal incentives. We formalize and solve this problem by maximizing social welfare under a budget constraint. The personalized incentives depend on the alternatives available to each individual and on her preferences. A polynomial time approximation algorithm computes a policy within few seconds. We analytically prove that it is boundedly close to the optimum. We efficiently calculate the curve of social welfare achievable for each value of budget within a given range. This curve can be useful for the regulator to decide the appropriate amount of budget to invest. We extend our formulation to enforcement, taxation and non-personalized-incentive policies. We analytically show that our personalized-incentive policy is also optimal within this class of policies and construct close-to-optimal enforcement and proportional tax-subsidy policies. We then compare analytically and numerically our policy with other state-of-the-art policies. Finally, we simulate a large-scale application to mode choice to reduce CO₂ emissions.

KEYWORDS

Personalized incentives; Knapsack problem; Tax policy; CO₂ emissions; Modal shift

JEL CLASSIFICATION

C61, H2, Q58, R41

This is an Accepted Manuscript of an article published by Taylor & Francis in *Transportmetrica A: Transport Science* on 23 Nov 2023, available at: <https://doi.org/10.1080/23249935.2023.2284353>.

1. Introduction

Taxes and subsidies in transportation are often perceived by the population as unfair, since they neglect the alternatives actually available to each individual and the individual preferences.¹ On the other hand, with the increase in information available to governments (Clarke and Margetts 2014), economic policies can be improved to consider the peculiarities of each individual. Customized policies could be used to align the individual cost with the social cost in the individuals’ decisions, without penaliz-

CONTACT Lucas Javaudin. Email: lucas.javaudin@cyu.fr

¹A preliminary version of this work was presented at the 2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC), and a conference paper appears in the proceedings of that conference (Javaudin, Araldo, and de Palma 2022).

ing anyone. We do not discuss here the acceptability and the legal dimension of such policies (which should be debated in the political arena).

We propose a policy of personalized incentives in a framework where individuals choose between multiple alternatives (or options). A benevolent regulator has a limited budget that he can use to propose monetary incentives, with the goal to induce individuals to change their choice toward socially-better ones. This policy is fair since no individual increases or decreases her utility. This is a clear advantage over road pricing which usually decreases the utility of some individuals.

Consider a regulator aiming to induce car buyers to choose more environmentally-friendly car models. Suppose that two buyers, A and B, both consider buying a car with high CO₂ emissions and suppose that buyer A (resp. buyer B) can be convinced to buy instead a car with low CO₂ emissions if she gets a discount of \$2000 (resp. of \$5000). With the policy of personalized incentives envisaged in this paper, the regulator could give \$2000 (resp. \$5000) to buyer A (resp. to buyer B) if she accepts to buy the low-emission car. In this simple example, the regulator could convince the two buyers to choose the low-emission car for only \$7000 (\$2000 to buyer A and \$5000 to buyer B), while, with a non-personalized subsidy policy, the regulator would need to spend at least \$10 000 to convince both buyers (each buyer receives \$5000). Hence, a personalized-incentive policy allows to reduce the average CO₂ emissions by the same amount than a non-personalized policy, while spending less.

We define the optimal personalized-incentive policy as the allocation of incentives that maximizes social welfare (defined as the reduction of CO₂ emissions in the example above), for a given budget. With two cars and two buyers, the optimal policy is easy to compute by simple enumeration. However, the problem is combinatorial so, with a large number of heterogeneous buyers and a large number of car models to choose from, we need more sophisticated methods.

We formalize the problem of finding a personalized-incentive policy maximizing social welfare under the regulator's budget constraint and show that it reduces to the well-known Multiple-Choice Knapsack Problem (MCKP – see Section 3). The MCKP consists in packing 'items' of different 'classes' into a knapsack of a certain 'capacity'. We show that the MCKP provides a natural formalization of an optimal incentive policy, where a class is an individual, an item is an alternative and the knapsack capacity is the budget of incentives. To approximate the optimal policy in polynomial time, we adapt a greedy algorithm from the Operations Research literature and we analyse some of its analytical (e.g., suboptimality gap bound) and economic (e.g., diminishing returns) properties (Section 4).

We then frame personalized-incentive policies into a larger family of demand management policies, including enforcement, tax and non-personalized-incentives (Section 5). These policies aim to maximize social welfare subject to a disutility constraint, where the disutility is the total loss of surplus for both the regulator and the individuals. We find that personalized-incentive policies are optimal within this larger family of policies. Moreover, they are 'fair', since they guarantee that the utility of each individual remains unchanged, and thus no one is penalized. We also compute a theoretical lower bound on the gap between state-of-the-art incentive policies, which are not personalized, and our personalized-incentive policy. Furthermore, we show that our greedy algorithm can not only construct incentive policies, but also enforcement and proportional tax-subsidy policies. We show that also in this case, the social welfare is boundedly close to the theoretical optimum. While in most of the paper we assume that the regulator knows exactly the preferences of each individual, we also study the case of imperfect information (Section 6).

Using data from the French census, we evaluate the CO₂ reduction achieved via the policy computed with our algorithm in a large-scale use-case, where individuals are incentivized to shift toward greener transportation modes for their commute to work, at the scale of a French department (Section 7). The results confirm the theoretical findings, showing in particular that our personalized incentives achieve the same CO₂ reduction as flat subsidies, but with a considerably smaller amount of incentives spent. Our code is available as open source.²

Even though the case study is about modal shift, the proposed methods can be applied in various contexts. For example, consider the marketing department of a large firm selling mutually exclusive goods. To increase the profits of the firm, the marketing department could use its budget to propose price discounts to some consumers in order to convince them to shift to goods with higher margins. Another potential example of application is in the telecommunications management context. In recent years, governments are planning to subsidize local organizations to improve the access of rural population to the Internet (France alone will spend 3 billions euros in 10 years, Arcep 2021). With our methods, governments could allocate optimally these subsidies.

2. Related Work

We first discuss the literature on incentive policies (Section 2.1), in particular in transportation, our main application domain. We then review the applications of the Multiple Choice Knapsack Problem, on which our optimization is based, in economics and transportation (Section 2.2) and, finally, in other domains (Section 2.3).

2.1. *Incentive Policies in Transportation*

Earlier studies of welfare analysis in a discrete-choice framework have been conducted by Small and Rosen (1981) and Anderson, de Palma, and Thisse (1992). De Borger (2001) studies the optimal taxation in a discrete-choice framework with externalities. His model is close to ours but he does not consider incentive policies. Some papers conduct an empirical study of an incentive policy in the transportation context (e.g., Merugu, Prabhakar, and Rama 2009, Etema, Knockaert, and Verhoef 2010, Knockaert et al. 2012, Yue et al. 2015, Hu, Chiu, and Zhu 2015, Li, Chen, and Tian 2021) but they do not carry out a theoretical study of the optimal policy.

Incentives are a promising tool for policy makers to trigger a transition toward greener transportation. Zhu et al. (2020) develop an approach to offer personalized incentives to a single individual to modify its travel behavior (e.g., mode, departure time or route chosen). We are interested instead in calculating optimal incentives for a large plethora of individuals. A vast literature exists on time-varying incentives and/or surcharges to shift departure times, in order to reduce congestion. To this aim, Sun et al. (2020) adopt a bottleneck model of a road segment. Tang et al. (2020) propose an optimization model to calculate transit surcharges and incentives, during peak and off-peak, respectively, to avoid over-crowding. In the two aforementioned works, the incentives are not personalized, in that they do not depend on the individual’s profile and all the individuals go from the same origin to the same destination. We instead consider a large set of individuals, each with a different set of alternatives, resulting in different contributions to the social welfare and individual perceived utility. Our in-

²<https://github.com/LucasJavaudin/individualized-incentives-algorithm>

centives are personalized, in that we encourage social-welfare maximizing alternatives with an incentive that compensate for the reduction in individual utility loss, which changes from an individual to another.

We now discuss three articles proposing a methodology to compute a personalized-incentive policy in large-scale scenario.

Araldo et al. (2019) devise Tripod, a simulation-based optimization method to calculate incentives to encourage energy efficient transportation alternatives. The incentives proposed depend on the alternatives available to the individuals but not on their preferences. Indeed, the system computes a unique ‘Token Energy Efficiency’ (TEE) value, and computes the incentive for each alternative by simply multiplying the TEE by the estimated energy savings achieved with that alternative. Such approach is pertinent when the regulator has no information on the individual preferences. However, when partial or perfect information is available, our approach is able to achieve a larger social welfare with the same incentives spent. We show these findings numerically in Section 7.6.

Ghafelebashi, Razaviyayn, and Dessouky (2023) induce drivers to select different routes to reduce congestion. Their findings indicate that providing an average incentive of \$2 leads to a 0.6% reduction in travel time, with 3% of drivers rewarded. In their approach, the incentive amount is limited to some arbitrary values, while in our approach, the incentive amount is computed in a preliminary step, using individual-level preferences.

Finally, Xiong et al. (2020) explore the effectiveness of incentives to influence not only mode choice but also departure time, route choice, and driving style. They find that the impact of incentives is more significant when compared to cases where incentives focus solely on altering route choice. In contrast, our approach is grounded in econometric principles as it employs behavioral discrete-choice analysis, a methodology commonly used in major transportation planning models.

2.2. *Multiple Choice Knapsack Problem in Economics and Transportation*

The Multiple Choice Knapsack Problem (MCKP – Kellerer et al. 2004, chap. 11) can be used to model a decision maker willing to optimally invest a limited budget in order to maximize an objective function. The possible investments options are divided into separate groups, and the decision maker has to choose at most one option for each group.

We now review the few examples of applications of MCKP in Economics and Transportation. Zhong and Young (2010) study the decision of a transportation planner willing to select a subset of candidate projects for funding. They do not propose any resolution algorithm and solve the problem in an exact way using a solver. Later, Colomi et al. (2017) use MCKP as a subroutine of a more general multi-criteria project-selection problem. Since the problem is NP hard (Kellerer et al. 2004, chap. 11), the aforementioned exact approach would require an unfeasibly large computation time in the large-scale applications we target. For this reason, we resort instead to a polynomial time approximation algorithm. Zoltners, Sinha, and Chong (1979, Sec. 6 and 7) use MCKP as a subroutine for a problem where a sales representative with a finite time-budget has to optimally allocate a call frequency to each accounts. They solve such a subroutine with an algorithm similar in spirit to our Algorithm 1, but in a more complicated setting, due to iterating decisions over multiple time-slots.

2.3. *Multiple Choice Knapsack Problem in Computer Science and other Domains*

MKCP is widely adopted in the Computer Science community, where a certain resource must be allocated among different entities. In the work of Cao, Brahma, and Varshney (2015), a central information aggregator receives information from several selfish sensors, which can transmit it at several precision levels: the more the precision level the higher the sensor cost in terms of energy. The aggregator needs to select one precision level (or none) per sensor and compensates the corresponding loss of energy of each sensor via payments. In Fielder et al. (2016), a manager of an information system invests in security controls. Per each control, it has to select a certain ‘level’: the higher the level, the higher the protection of the organization, but also the higher the investment. Araldo, Di Stefano, and Di Stefano (2020) allocate computational resources among different service providers; the owner of the resources selects one configuration for each of them in order to increase the overall system utility. To this aim, they use Multi-Dimensional Multiple-Node MCKP. Mohammadivojdan and Geunes (2018) solve the problem of a seller, who needs to decide the amount of product to buy from several providers, each proposing a different pricing scheme, in order to maximize its overall utility.

2.4. *Position with respect to the Related Work*

To the best of our knowledge, we are the first to formalize the problem of computing optimal personalized-incentives with MCKP. By finding the assumptions that enable such a formalization, we show in this paper that MCKP describes naturally such a problem, since it manages to represent the different alternatives of each individual. The adoption of MCKP also allows us to devise an efficient algorithm for large-scale applications, adapting existing solutions from Operations Research.

3. Framework and Personalized-Incentive Policy

In this section, we first formalize the model studied and present the underlying assumptions (Section 3.1). We also characterize the personalized-incentive policy that will be studied throughout this paper (Section 3.2). We then present the Maximum Social Welfare Problem, which consists in finding the optimal incentive policy under a budget constraint (Section 3.3).

The notations used throughout this paper are summarized in Table 1. All the proofs are relegated to Appendix A.

3.1. *Model and Assumptions*

We consider a population $\mathcal{I} \equiv \{1, \dots, m\}$ of m individuals. Each individual $i \in \mathcal{I}$ chooses an alternative j among an *individual-specific choice-set* \mathcal{N}_i . For example, we can consider individuals choosing a mode of transportation to commute to their work. In this case, the choice set could be $\mathcal{N}_i = \{\text{car, walk, bike, public transit}\}$. The choice set can be individual-specific so that if individual i owns a car but individual i' does not, we could have $\mathcal{N}_i = \{\text{car, walk, bike, public transit}\}$ and $\mathcal{N}_{i'} = \{\text{walk, bike, public transit}\}$. The mode-choice example is studied extensively in Section 7. As another example, \mathcal{I} could be a set of individuals purchasing a car. In this

Table 1. Notation used throughout the paper.

i	Index to denote an individual
j	Index to denote an alternative
$[i, j]$	Alternative j of individual i
\mathcal{I}, m	Set of individuals and number of individuals
\mathcal{N}_i	Set of alternatives available to individual i
$V_{i,j}$	Intrinsic utility of individual i when choosing alternative j
$b_{i,j}$	Social indicator of alternative j of individual i
$z_{i,j}$	Monetary transfer received (or paid) by individual i when choosing alternative j
\mathbf{z}	General policy (set $\{z_{i,j}\}_{i,j}$ of monetary transfers)
$y_{i,j}$	Personalized incentive proposed to individual i , conditional on choosing alternative j
\mathbf{y}	Personalized-incentive policy (set $\{y_{i,j}\}_{i,j}$ of incentives)
\mathcal{Y}	Set of all personalized-incentive policies
$U_{i,j}(\mathbf{z})$	Utility of individual i when choosing alternative j , given policy \mathbf{z} (1)
$j_i^*(\mathbf{z})$	Alternative chosen by individual i , given policy \mathbf{z}
$j_i^*(0)$	Default alternative, i.e., alternative chosen by individual i , in the absence of policy (5)
$w_{i,j}$	Weight of alternative j of individual i , equation (8)
$e_{i,j}$	Efficiency of alternative j of individual i (Definition 4.1)
\bar{Q}	Maximum budget available to the regulator
\hat{Q}	Budget actually spent by the policy computed by Algorithm 1
$B^*(Y)$	Maximum social welfare reachable with a personalized-incentive policy, with a total incentive expenditure of Y
$B(Y)$	Social welfare obtained with the personalized-incentive policy produced by our algorithm, with a total incentive expenditure of Y
$B(\mathbf{z})$	Social welfare achieved with a policy \mathbf{z} , equation (2)
$Y(\mathbf{z})$	Expenses of the regulator for a policy \mathbf{z} , equation (3)
$\Delta U(\mathbf{z})$	Total variation in individual utility of a policy \mathbf{z} , equation (18)
$\delta(\mathbf{z})$	Disutility of policy \mathbf{z} , equation (19)
$\mathcal{R}_i \subseteq \mathcal{N}_i, r_i = \mathcal{R}_i $	Set of LP-extremes alternatives of individual i , and its cardinality
$\tilde{b}_{i,j}$	Incremental social indicator of the alternative j of individual i
$\tilde{w}_{i,j}$	Incremental incentive of the alternative j of individual i
$\tilde{e}_{i,j}$	Incremental efficiency of the alternative j of individual i
$\tilde{e}_{s,t}$	Incremental efficiency of the split item (Algorithm 1)
k	Iteration index of Algorithm 1
$e^{[k]}, \tilde{e}^{[k]}$	Overall and incremental efficiency of Algorithm 1 at iteration k (Definition 4.6)
τ	Tax level (Section 5.2.2)
A_i	Baseline social-indicator of individual i (Section 5.2.2)

case, the set of alternatives \mathcal{N}_i of individual i would include the models of cars available in the market.

Let $z_{i,j} \in \mathbb{R}$ be a monetary transfer, from the regulator to individual i , induced when she chooses alternative j . This monetary transfer can be an incentive, if positive, or a tax, if negative. Any *policy* can thus be described by a set of monetary transfers proposed to all the individuals for any of their alternatives, which we compactly denote with $\mathbf{z} \equiv \{z_{i,j}\}_{i,j}$.

A policy influences the individual choice since the proposed monetary transfers change her utilities.

The *utility* $U_{i,j}(\mathbf{z})$ of individual i when choosing alternative $j \in \mathcal{N}_i$ is given by

$$U_{i,j}(\mathbf{z}) = V_{i,j} + z_{i,j}, \quad (1)$$

where $V_{i,j} \in \mathbb{R}$ is the intrinsic utility (in the absence of policy). We implicitly assumed that utility is quasi-linear with respect to income, which means that both $V_{i,j}$ and $z_{i,j}$ are expressed in the same unit as income and that $z_{i,j}$ has an additive effect on utility, hence equation (1).

Given a policy \mathbf{z} , each individual i chooses an alternative $j_i^*(\mathbf{z})$ which maximizes her utility, i.e., $j_i^*(\mathbf{z}) \in \arg \max_j U_{i,j}(\mathbf{z})$.

We consider a regulator aiming to maximize a social welfare indicator, whose value depends on the individuals' choices. More formally, each alternative j of individual i is characterized by a *social indicator* $b_{i,j} \in \mathbb{R}$. In the mode-choice example, the social indicator could be the opposite of CO₂ emissions induced by the commutes.

The goal of the regulator is to find a policy \mathbf{z} which maximizes the global social indicator, or *social welfare* indicator, defined by

$$B(\mathbf{z}) \equiv \sum_{i=1}^m b_{i,j_i^*(\mathbf{z})}, \quad (2)$$

i.e., the sum of the social indicators of the alternatives chosen by the individuals. Intuitively, a policy \mathbf{z} which maximizes welfare could be such that $z_{i,j} = 0$ if $j \in \arg \max_{j'} b_{i,j'}$, and $z_{i,j} = -\infty$ otherwise, for any individual i and any alternative j , which is equivalent to a ban of all alternatives that do not maximize the social indicator for each individual. However, in practice, the regulator is affected by some political constraints and such extreme policy is not feasible.

Definition 3.1 (Expenses). For any policy \mathbf{z} , we define the expenses $Y(\mathbf{z})$ of the regulator (or his revenues $-Y(\mathbf{z})$) as

$$Y(\mathbf{z}) \equiv \sum_{i=1}^m z_{i,j_i^*(\mathbf{z})}, \quad (3)$$

i.e., the sum of the monetary transfers paid or received for the alternative $j_i^*(\mathbf{z})$ chosen by each individual i .

The following assumptions are made. First, we assume that individuals cannot affect each other's intrinsic utilities and social indicators.

Assumption 3.2 (Independent intrinsic utilities). Given a policy \mathbf{z} , for each individual i and each alternative $j \in \mathcal{N}_i$, the intrinsic utility $V_{i,j}$ is independent of the

alternative $j_i^*(\mathbf{z})$ chosen by any other individual $i' \neq i$.

Assumption 3.3 (Independent social indicators). Given a policy \mathbf{z} , for each individual i and each alternative $j \in \mathcal{N}_i$, the social indicator $b_{i,j}$ is independent of the alternative $j_i^*(\mathbf{z})$ chosen by any other individual $i' \neq i$.

Note that Assumptions 3.2 and 3.3 imply that our policy has negligible effect on congestion and crowding. This holds in many practical situations. For instance, in the numerical scenario on transport mode choice (Section 7), we achieve relevant social welfare improvement (CO2 reduction), while inducing only few individuals to change their modes, leading to a negligible impact on congestion.

We further assume that the utilities and social indicators are known to the regulator.

Assumption 3.4 (Perfect information). The regulator has perfect information: it knows exactly the intrinsic utilities $\{V_{i,j}\}_{i,j}$ and social indicators $\{b_{i,j}\}_{i,j}$ of all the alternatives, for all the individuals.

Assumption 3.4 enables us to build our theoretical framework. However, it will be relaxed in Sections 6 and 7.7 by assuming that the regulator has only imperfect knowledge on the intrinsic utilities. It is also conceivable that the regulator could elicit user preferences through a substantial number of repeated observations, mirroring the common practices seen in the marketing activities of private companies today. Such approaches are widely accepted in the private sector, making a case for extending similar strategies to a regulator aiming to promote social welfare. While this could introduce significant concerns regarding public acceptance, our paper can shed light on the potential societal gains attainable by allowing the regulator to gather substantial amounts of personalized information.

With no loss of generality, we rule out identical alternatives.

Assumption 3.5 (No identical alternatives). For any individual i , there are no identical alternatives $j, j' \in \mathcal{N}_i$, i.e., such that $V_{i,j} = V_{i,j'}$ and $b_{i,j} = b_{i,j'}$.

We need to characterize more precisely the behaviour of individuals when multiple alternatives maximize their utility.

Assumption 3.6 (Tie-breaking rule). For any policy \mathbf{z} , if the set $\arg \max_j U_{i,j}(\mathbf{z})$ contains more than one alternative, individual i chooses the alternative j' with the largest social indicator, i.e.,

$$j_i^*(\mathbf{z}) = \arg \max_{j' \in \arg \max_j U_{i,j}(\mathbf{z})} b_{i,j'}. \quad (4)$$

The previous assumption is merely a technical assumption that could be relaxed by proposing incentives infinitesimally larger to ensure that the set $\arg \max_j U_{i,j}(\mathbf{z})$ is always a singleton.

The alternative chosen by each individual i in the absence of policy (i.e., where $z_{i,j} = 0, \forall i, j$) is called *default alternative*, and denoted $j_i^*(0)$. Under Assumption 3.6, the default alternative is given by

$$j_i^*(0) \equiv \arg \max_{j' \in \arg \max_j V_{i,j}} b_{i,j'}. \quad (5)$$

3.2. Personalized-Incentive Policies

We assume that the space of policies available to the regulator is limited to policies such that $z_{i,j} \geq 0$, for each alternative j and individual i . In other words, the regulator never taxes alternatives, for some political reasons. Note that we allow the regulator to give different monetary transfers to different individuals for the same alternative (e.g., some individuals might receive \$2 for commuting by foot, while others may only receive \$1). Hence, this space of policies is referred to as the set of *personalized-incentive policies*, denoted \mathcal{Y} . To distinguish personalized-incentive policies from more general policies, we denote them with $\mathbf{y} = \{y_{i,j}\}_{i,j}$, where $y_{i,j}$ is the incentive given to individual i , conditional on her choosing alternative j , and thus,

$$\mathcal{Y} = \{\mathbf{y} = \{y_{i,j}\}_{i,j} : y_{i,j} \geq 0, \forall i, j\}.$$

The incentive $y_{i,j}$ reduces the budget of the regulator only if individual i chooses alternative j . Therefore, if the regulator wants to spend at most a budget Q , the budget constraint can be written as

$$Y(\mathbf{y}) = \sum_{i=1}^m y_{i,j_i^*(\mathbf{y})} \leq Q,$$

where $j_i^*(\mathbf{y})$ is the alternative chosen by individual i under the personalized-incentive policy \mathbf{y} .

In the rest of this subsection, we characterize more precisely the set of policies we consider, discarding ‘inefficient’ policies.

Proposition 3.7. *The regulator can induce any individual $i \in \mathcal{I}$ to shift from her default alternative $j_i^*(0)$ to any alternative $j \in \mathcal{N}_i$, with a higher social indicator (i.e., $b_{i,j} > b_{i,j_i^*(0)}$), by proposing the following incentives*

$$\begin{cases} y_{i,j} &= V_{i,j_i^*(0)} - V_{i,j}, \\ y_{i,j'} &= 0, \end{cases} \quad \text{for any other alternative } j' \neq j \quad (6)$$

Additionally, $y_{i,j}$, defined above, is the minimum incentive required to induce individual i to shift to alternative j .

Such a proposition tells us that it suffices to incentivize only one alternative per individual and no more than that. Therefore, we can limit the space of the studied personalized-incentive policies as in the following assumption, with no loss of generality.

Assumption 3.8. We only study personalized-incentive policies that propose incentives in the form of (6), i.e., only one alternative j per individual i is incentivized, with an incentive equal to $y_{i,j} = V_{i,j_i^*(0)} - V_{i,j}$.

If an individual i is given an incentive $y_{i,j}$, given by (6), to shift to alternative j , then her utility $U_{i,j}(\mathbf{y})$ remains unchanged since, from equation (1),

$$U_{i,j}(\mathbf{y}) = V_{i,j} + \underbrace{V_{i,j_i^*(0)} - V_{i,j}}_{y_{i,j}} = V_{i,j_i^*(0)}.$$

In other words, the incentive amount is such that the utility of the individual does not change. In this sense, our approach is fully equitable.

3.3. Maximum Social Welfare Problem

We can now formally define the optimization problem of the regulator, who chooses the personalized-incentive policy $\mathbf{y} = \{y_{i,j}\}_{i,j}$ which maximizes social welfare under his budget constraint. We refer to this problem as *Maximum Social Welfare Problem*:

$$\begin{cases} \max_{\mathbf{y} \in \mathcal{Y}} & B(\mathbf{y}) \\ \text{s.t.} & Y(\mathbf{y}) \leq Q \\ & y_{i,j} \geq 0, \quad \forall i \in \mathcal{I}, j \in \mathcal{N}_i \end{cases} . \quad (7)$$

Definition 3.9 (Optimal personalized-incentive policy). An optimal personalized-incentive policy \mathbf{y} , for a budget Q , is a solution of Problem 7.

We denote with $\{j_1, \dots, j_m\}$ a *chosen-alternative set*, where j_i denotes the alternative chosen by individual $i \in \mathcal{I}$. According to Proposition 3.7, the regulator can induce any chosen-alternative set $\{j_1, \dots, j_m\}$ by proposing to any individual i the incentives $y_{i,j_i} = V_{i,j_i^*}(0) - V_{i,j_i}$ and $y_{i,j} = 0$, for any $j \neq j_i$. Thanks to the same proposition, the regulator cannot induce this set of alternatives by spending less. Therefore, the optimization problem of the regulator (7) amounts to finding the chosen-alternative set $\{j_1, \dots, j_m\}$ which maximizes social welfare $\sum_{i=1}^m b_{i,j_i}$, subject to the constraint that the corresponding spendings $\sum_{i=1}^m y_{i,j_i}$ must not exceed the budget Q .

Such a problem can be expressed as an Integer Linear Program (ILP). In order to do so, we introduce a *weight* $w_{i,j}$ for any alternative $j \in \mathcal{N}_i$ of individual i . The weight is defined as the incentive amount that would be proposed to individual i if the regulator were to induce her to choose alternative j , which is, according to Proposition 3.7,

$$w_{i,j} \equiv V_{i,j_i^*}(0) - V_{i,j}, \quad \forall i, j. \quad (8)$$

Note that $w_{i,j}$ is a fixed value that we used to compute the optimal policy, while $y_{i,j}$ represents the incentive amount chosen by the regulator. The personalized-incentive policy is such that $y_{i,j} = w_{i,j}$, if individual i is induced to choose alternative j , and $y_{i,j} = 0$ otherwise.

We introduce the binary decision variable $x_{i,j}$ that is equal to 1 if the regulator wants to make individual i choose alternative j , and that is equal to 0 otherwise, with the natural constraint that $\sum_{j \in \mathcal{N}_i} x_{i,j} = 1$ (only one alternative is chosen). The Maximum Social Welfare problem (7) can be written as

$$\begin{cases} \max_{\{x_{i,j}\}_{i,j}} & \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{N}_i} b_{i,j} x_{i,j} \\ \text{s.t.} & \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{N}_i} w_{i,j} x_{i,j} \leq Q \\ & \sum_{j \in \mathcal{N}_i} x_{i,j} = 1, & i \in \mathcal{I} \\ & x_{i,j} \in \{0, 1\}, & i \in \mathcal{I}, j \in \mathcal{N}_i \\ & w_{i,j} = V_{i,j_i^*}(0) - V_{i,j}, & i \in \mathcal{I}, j \in \mathcal{N}_i \end{cases} , \quad (9)$$

which is a Multiple-Choice Knapsack Problem (MCKP) with weights $w_{i,j}$ and profits $b_{i,j}$ (Kellerer et al. 2004, chap. 11).

Observe that the solution $\{x_{i,j}\}_{i,j}$ of problem (9) corresponds to the personalized-incentive policy \mathbf{y} , solution of (7), where

$$y_{i,j} = x_{i,j} \cdot w_{i,j}, \quad \forall i, j.$$

We denote by $B^*(Q)$ the maximum of the social welfare, solution of problem (9).

3.4. *Maximum Social Welfare Curve Problem*

Suppose now that the regulator is endowed with a maximum budget Q and that he can spend any budget in the interval $Y \in [0, Q]$. To decide the exact amount of budget to spend, it is useful to obtain the *Maximum Social Welfare Curve* \mathcal{C}_Q^* , representing the maximum social welfare reachable, $B^*(Y)$, for any budget $Y \in [0, Q]$, i.e.

$$\mathcal{C}_Q^* = \{(Y, B^*(Y)) \mid Y \in [0, Q]\}. \quad (10)$$

The *Maximum Social Welfare Curve Problem* consists in finding the curve \mathcal{C}_Q^* , for a given maximum budget Q . It is easy to show that it is monotone non-decreasing (the larger the budget spent, the larger the social welfare reached). Observe that, although a maximum budget Q is available, the regulator may not want to indiscriminately spend it all, but may choose the actual budget to invest in incentives, based on several criteria. For instance, the regulator may use the above curve to find the minimum budget needed to reach a certain social-welfare target. Moreover, in many practical cases, the social welfare is converted into money metric. The coefficient of conversion is usually fixed based on political considerations. For instance, in our numerical results (Section 7), we convert CO₂ emission reduction into money, using the cost of 100 euros per ton of CO₂. After converting social welfare in money metric, the regulator may choose to invest an incentive budget such that the gain of social welfare equals the incentive spent. Such a value can be found on the Maximum Social Welfare Curve.

4. Approximation Algorithm

Kellerer et al. (2004) shows that the MCKP problem, and thus the Maximum Social Welfare problem (9), is NP-hard. Therefore, for large instances of such problems, finding the optimal solution is unfeasible and we need to resort to heuristics. We provide in this section a polynomial time algorithm based on greedy algorithms from the Operations Research literature, which gives us solutions boundedly close to the optimum. We then discuss some relevant properties of such an algorithm, e.g., its diminishing returns behaviour and the fact that it is an anytime algorithm. In the following subsection, we introduce some preliminary mathematical concepts.

4.1. *Preliminary Steps*

Before presenting the proposed algorithm, we need to ‘clean’ the input of the problem, removing some irrelevant alternatives from the set \mathcal{N}_i of the alternatives of any individual i . Irrelevant alternatives are the ones that do not provide enough social indicator compared to the incentive amount needed to induce them. We call *LP-extremes* the alternatives remaining after the cleaning, and we denote them with $\mathcal{R}_i \subseteq \mathcal{N}_i$. The

process of constructing the set \mathcal{R}_i is called *concavization* and is described in detail in Kellerer et al. (2004, Section 11.2.1). See also Javaudin, Araldo, and de Palma (2022, Section III. A.) for an intuitive visualization of the concavization process.

We follow the Operations Research literature in the slight abuse of notation of denoting with $w_{i,j}$ the incentive to be provided to the j -th alternative in \mathcal{R}_i , where this is not ambiguous. With no loss of generality, we can assume the ordering

$$w_{i,1} < w_{i,2} < \dots < w_{i,r_i} \quad (11)$$

in the set \mathcal{R}_i , where r_i is its cardinality. By definition, the default alternative requires no incentive ($w_{i,j_i^*(0)} = 0$), and thus it is the first alternative in the set \mathcal{R}_i and $w_{i,1} = 0$.

Definition 4.1 (Efficiency and incremental efficiency). We define the *efficiency* of an alternative j of individual i as

$$e_{i,j} \equiv (b_{i,j} - b_{i,j_i^*(0)})/w_{i,j},$$

i.e., the gain in social indicator that we can gain via a unit of incentive allocated to that alternative. We define the incremental social indicator $\tilde{b}_{i,j}$ and the incremental incentive $\tilde{w}_{i,j}$ required for each alternative $j \in \mathcal{R}_i$ as

$$\begin{aligned} \tilde{b}_{i,j} &\equiv b_{i,j} - b_{i,j-1} \\ \tilde{w}_{i,j} &\equiv w_{i,j} - w_{i,j-1} \end{aligned}, \quad j = 2, \dots, r_i. \quad (12)$$

The *incremental efficiency* is then defined as

$$\tilde{e}_{i,j} \equiv \tilde{b}_{i,j}/\tilde{w}_{i,j}. \quad (13)$$

The incremental efficiency $\tilde{e}_{i,j}$ represents the increase in social welfare for each monetary unit spent, when individual i shifts from alternative $j - 1$ to alternative j .

4.2. Greedy Algorithm

We want to find a curve $\mathcal{C}_Q = \{(Y, B(Y)) \mid Y \in [0, Q]\}$ that approximates the Maximum Social Welfare Curve \mathcal{C}_Q^* (10), i.e., such that $B(Y)$ is close to $B^*(Y)$ for any value of budget $Y \geq 0$.

Very efficient algorithms, like the Dyer-Zemel algorithm (Kellerer et al. 2004, Section 11.2.1) are known to solve problem (9), i.e., to approximate the maximum social welfare for a fixed single value of budget Q . However, to apply them to the Maximum Social Welfare Curve problem, in which we want to find the maximum social welfare for a range of budget values $Y \in [0, Q]$, instead of just one, we would have to run those algorithms from scratch for every single value of budget. For this reason, we build our solutions upon a simpler greedy algorithm (Kellerer et al. 2004, Figure 11.2), which is less efficient to solve the Maximum Social Welfare problem (although still polynomial in time complexity), but easily extendable to also solve the Maximum Social Welfare Curve problem. The other advantage deriving from such choice is that this greedy algorithm has interesting properties that increase its practical application and economic interpretability, as discussed in Section 4.3.

The pseudocode of the algorithm is in Algorithm 1. The notation $[i, j]$ stands for ‘ j -th alternative of individual i ’. The idea of the algorithm is simple. First, the algorithm

finds all the LP-extremes alternatives and sorts them by order of decreasing incremental efficiency. Then, at each iteration, the next pair of individual and alternative $[i', j']$ with the highest incremental efficiency is picked (line 8). The alternative induced to i' is set to j' (line 10) and the budget is reduced by the amount of the incremental weight (equation (14)). An additional piece of the approximation of the social welfare curve is computed (equation (16)). The algorithm stops when the maximum budget Q is depleted and it returns a policy \mathbf{y} , which is such that any individual i , for whom the algorithm selected an alternative $j \in \mathcal{R}_i$, effectively chooses this alternative j .

Algorithm 1: Greedy algorithm for the Maximum Social Welfare and Maximum Social Welfare Curve problems.

Input : Social indicators $\{b_{i,j}\}_{i,j}$, intrinsic utilities $\{V_{i,j}\}_{i,j}$, budget Q

- 1 Iteration index $k := 0$
- 2 $Y^{[k]} := 0$, Total incentive allocated so far.
- 3 $B^{[k]} := 0$, Social welfare obtained in the current allocation.
- 4 Compute the ordered set \mathcal{R}_i of LP-extremes of each individual i .
- 5 Sort all the alternatives $[i, j]$ according to decreasing incremental efficiency $\tilde{e}_{i,j}$ and put them in a set \mathcal{R} .
- 6 Initialize the alternatives chosen by the individuals $\{x_{i,j}\}_{i,j}$ as follows

$$\begin{cases} x_{i,1} = 1, & \text{(default alternative)} \\ x_{i,j} = 0, & \text{for any alternative } j > 1 \end{cases}$$
- 7 **while** $\mathcal{R} \neq \emptyset$ and $Y^{[k]} \leq Q$ **do**
- 8 Take $[i', j']$, the next alternative with the highest incremental efficiency $\tilde{e}_{i',j'}$ from \mathcal{R} .
- 9 Add $[i', j']$ to the solution, i.e.:

$$\begin{aligned} \mathcal{R} &:= \mathcal{R} \setminus \{[i', j']\}, \\ Y^{[k+1]} &:= Y^{[k]} + \tilde{w}_{i',j'} & (14) \\ \tilde{e}^{[k]} &:= \tilde{e}_{i',j'} & (15) \\ B(Y) &:= B^{[k]}, & \forall Y \in [Y^{[k]}, Y^{[k+1]}) & (16) \\ B^{[k+1]} &:= B^{[k]} + \tilde{b}_{i',j'} \\ k &:= k + 1 \end{aligned}$$
- 10 Update the selected alternative for individual i' , i.e.,

$$\begin{cases} x_{i',j'} = 1, \\ x_{i',j} = 0, & \text{for any other alternative } j \neq j' \end{cases}$$
- 11 **end**

Output: Curve $\mathcal{C}_Q = \{(Y, B(Y)) \mid Y \in [0, Q]\}$
 Chosen alternatives $\{x_{i,j}\}_{i,j}$
 Incentive policy $\mathbf{y} = \{y_{i,j}\}_{i,j}$, where $y_{i,j} = x_{i,j} \cdot w_{i,j}$
 Split item $[s, t] := [i', j']$
 Incremental efficiency of the split item $\tilde{e}_{s,t}$
 Budget actually used $\tilde{Q} := Y^{[k-1]}$

Observe that the curve \mathcal{C}_Q given as output by the algorithm is an approximation of the solution \mathcal{C}_Q^* of the Maximum Social Welfare Curve Problem (Section 3.4). Moreover, given any maximum budget Q , the algorithm returns an approximation $B(Q)$ to the solution $B^*(Q)$ of the Maximum Social Welfare Problem (9). Note that, in order to achieve $B(Q)$, the policy issued by the algorithm does not spend the entire maximum budget Q , but only $\tilde{Q} \leq Q$.

The algorithm also gives as output the incremental efficiency of the ‘split item’, denoted with $\tilde{e}_{s,t}$, useful to compute the optimality gap of the algorithm (Theorem 4.2 below). The name *split item*, which we borrow from Kellerer et al. (2004), reminds of




Figure 1. Distance between the social welfare curve \mathcal{C}_Q computed by Algorithm 1, the maximum social welfare curve \mathcal{C}_Q^* (Section 3.4) and the upper bound of Theorem 4.2. The stars represent the incentive spent $Y^{[k]}$ and social welfare $B^{[k]} = B^*(Y^{[k]})$ at each iteration $k = 1, \dots, 8$ of the algorithm (line 9).

the fact that, when we allocate the budget Q , we add to the solution all the LP-extreme alternatives, in decreasing order of incremental efficiency, up to $[s, t]$. In other words, such alternative $[s, t]$ splits the set \mathcal{R} of all the LP-extremes in two parts: the first part consists of the alternatives we include in our solution, while we do not include the LP-extremes from the second part.

4.3. Useful Properties for Large-Scale Applications

The distance to the optimum, in terms of social welfare, is bounded from above.

Theorem 4.2 (Upper bound). *Let us run Algorithm 1 with budget Q , and let \tilde{Q} be the budget actually used and $\tilde{e}_{s,t}$ be the incremental efficiency of the split item. The social welfare $B(Q)$ we obtain is boundedly close to the social welfare $B^*(Q)$ of any optimal personalized-incentive policy (Definition 3.9). In particular,*

$$B^*(Q) - B(Q) \leq \tilde{e}_{s,t} \cdot (Q - \tilde{Q}). \quad (17)$$

Corollary 4.3. *The curve \mathcal{C}_Q obtained via Algorithm 1 is boundedly close to the Maximum Social Welfare Curve \mathcal{C}_Q^* from equation (10) and the gap is given by Theorem 4.2.*

The next corollary states that, for any budget Q , the curve \mathcal{C}_Q returned by Algorithm 1 and the Maximum Social Welfare Curve \mathcal{C}_Q^* ‘touch each other’. This ensures that the allocation computed by Algorithm 1 at every iteration is optimal. It is a direct consequence of Theorem 4.2.

Corollary 4.4. *The curve \mathcal{C}_Q obtained via Algorithm 1 and the Maximum Social Welfare Curve \mathcal{C}_Q^* from equation (10) are such that $B(Y^{[k]}) = B^*(Y^{[k]})$ in all the values $Y^{[k]}$, $k = 0, 1, \dots$*

Figure 1 illustrates this property. The continuous curve represents the Maximum Social Welfare Curve \mathcal{C}_Q^* and the dashed curve represents the curve \mathcal{C}_Q obtained via Algorithm 1. These two curves are step-functions because of the discreteness of the problem. From Corollary 4.4, the curve \mathcal{C}_Q intersects the curve \mathcal{C}_Q^* at each iteration of the algorithm (represented by the stars). The dotted curve represents the upper bound of \mathcal{C}_Q^* , computed from Theorem 4.2.

In the small-scale scenario of Figure 1, there are 10 individuals with 5 alternative each and it is thus possible to compute \mathcal{C}_Q^* using brute force (i.e., all the possible combinations of alternatives chosen are evaluated). Brute force requires comparing m^N policies, for N individuals with m alternatives each. Computing \mathcal{C}_Q^* is therefore unfeasible in realistic cases and we have to resort to an approximation algorithm.

Proposition 4.5. *The computational complexity of Algorithm 1 is $O(\sum_{i=1}^m |\mathcal{N}_i| \cdot \log |\mathcal{R}_i| + |\mathcal{R}| \cdot \log m)$, where m is the number of individuals, $|\mathcal{N}_i|$ is the number of alternatives of individual i , $|\mathcal{R}_i|$ is the number of LP-extremes of individual i and $|\mathcal{R}| \equiv \sum_{i=1}^m |\mathcal{R}_i|$.*

Note that, since the alternatives of each individual are independent of the others, the sets \mathcal{R}_i can be computed in parallel, thus reducing even further the computation time.

Two additional properties are useful when using the algorithm in large-scale scenarios: (i) the algorithm is “anytime”, i.e., if we stop it at any iteration, we get a valid allocation policy; (ii) we can build on a previously computed policy if new budget becomes available. These properties are detailed in Remarks III-B.4 and III-B.5 of Javaudin, Araldo, and de Palma (2022). Moreover, at any iteration we can compute how much the social welfare would increase if we continued until the end. If this increase is not relevant, we can decide to stop prematurely. We omit the details on the computation of this increase (see Araldo, Di Stefano, and Di Stefano 2020, Section 6.2.2 for details).

In order to describe the diminishing return property of Algorithm 1, we need the following definition.

Definition 4.6 (Incremental and overall efficiency). The *incremental efficiency* provided by the algorithm at iteration k is $\tilde{e}^{[k]}$, defined in equation (15). We define the *overall efficiency* of a personalized-incentive policy spending budget Y and achieving social welfare B as $e = B/Y$. We denote with $e^{[k]}$ the overall efficiency of the policy obtained by stopping Algorithm 1 at iteration k , i.e., $e^{[k]} = B^{[k]}/Y^{[k]}$.

The following proposition illustrates that, by spending more and more budget and allocating it as the algorithm dictates, we increase social welfare, but the marginal gain per unit of budget spent decreases.

Proposition 4.7 (Diminishing returns). *The incremental efficiency $\tilde{e}^{[k]}$ and the overall efficiency $e^{[k]}$ of the alternative added by Algorithm 1 at every iteration k are both monotonically non-increasing.*

In some cases, the regulator would be willing to maximize social welfare under the constraints that the *overall inverse efficiency* e^{-1} is below a certain target. For instance, in Section 7.5 the regulator does not want to spend more than 100 euros per ton of CO₂ saved, which is considered to be the carbon price. In such cases, it is useful to observe that, thanks to Proposition 4.7, $(e^{[k]})^{-1}$ is non-decreasing. Therefore, the regulator could run the algorithm and stop at the iteration where $(e^{[k]})^{-1}$ goes above the target inverse efficiency.

The properties presented in this section have shown that the proposed Algorithm 1 is computationally efficient and able to return an allocation providing a welfare close to the optimum. Moreover, it has some features that make its adoption easier in practical large-scale scenarios.

5. Comparison with Other Policies

So far, we have considered the case of a regulator endowed with a fixed budget and who is limited to giving subsidies (i.e., taxation is not an option, for legal or political reasons). We have shown that our personalized-incentive policy is boundedly close to the optimal policy in this case. In this section, to frame the policy into a more general set of feasible policies, we generalize the framework to include not only personalized incentives, but also enforcement policies, taxation, and non-personalized-incentive policies. In Section 5.1, we show that any optimal personalized-incentive policy (Definition 3.9) is optimal in this more general class of policies. In Section 5.2, we show that Algorithm 1 can be used to compute an enforcement policy and a proportional tax-subsidy policy, which are both boundedly close to the optimal general policy.

5.1. Optimality of Personalized-Incentives Policies among General Policies

We now consider a more general space of policies, including incentives, enforcement and taxation policies, and we define a criteria of optimality in this space. In order to do so, we need to define some new quantities.

The total loss in individual utility, of a policy \mathbf{z} , is

$$\Delta U(\mathbf{z}) \equiv \sum_{i=1}^m (V_{i,j_i^*(0)} - U_{i,j_i^*(\mathbf{z})}(\mathbf{z})). \quad (18)$$

With a taxation policy, i.e., a policy \mathbf{z} such that $z_{i,j} \leq 0$, $\forall i, j$, the loss in individual utility is non-negative, i.e., $\Delta U(\mathbf{z}) \geq 0$. With an incentive policy, i.e., a policy \mathbf{z} such that $z_{i,j} \geq 0$, $\forall i, j$, the loss in individual utility is non-positive, i.e., $\Delta U(\mathbf{z}) \leq 0$. In particular, in any personalized-incentive policy obeying to Proposition 3.7, the individuals are perfectly compensated for their loss in utility, and thus $\Delta U(\mathbf{z}) = 0$. Keeping everything else fixed, it is obvious that, the smaller ΔU , the better.

The disutility, or cost, of a policy is measured by combining the loss in individual utilities and the expenses for the regulator, defined in equation (3).

Definition 5.1 (Disutility). The disutility $\delta(\mathbf{z})$ of a policy \mathbf{z} is defined as the expenses of the regulator $Y(\mathbf{z})$ plus the total loss in individual utilities $\Delta U(\mathbf{z})$, i.e.,

$$\delta(\mathbf{z}) \equiv Y(\mathbf{z}) + \Delta U(\mathbf{z}). \quad (19)$$

The following proposition shows that the disutility of a policy is always non-negative, which means that it is not possible that both the individuals increase their utility and the regulator collects revenues. It also shows that, if two policies imply the same alternatives chosen, then they have the same disutility.

Proposition 5.2. *Every policy \mathbf{z} has a non-negative disutility that only depends on the alternative chosen by the individuals, rather than the actual incentive or taxation proposed. In particular:*

$$\delta(\mathbf{z}) = \sum_{i=1}^m (V_{i,j_i^*(0)} - V_{i,j_i^*(\mathbf{z})}) \geq 0$$

We can now define an optimal general policy, whose definition includes incentive, taxation and enforcement policies.

Definition 5.3 (Optimal General Policy). An optimal general policy \mathbf{z} with disutility threshold $Q \geq 0$ is the solution of the following problem:

$$\begin{cases} \max_{\mathbf{z}} & B(\mathbf{z}) \\ \text{s.t.} & \delta(\mathbf{z}) \leq Q \end{cases} \quad (20)$$

Note that, problem (20) is a generalization of (7). Indeed, we obtain the latter from the former by (i) constraining the policy to be a personalized-incentive policy, i.e., $z_{i,j} \geq 0, \forall i \in \mathcal{I}, j \in \mathcal{N}_i$ and (ii) imposing no change in individual utility, i.e., $\Delta U(\mathbf{z}) = 0$.

The two next propositions characterize the optimal general policy. The first one implies that finding an optimal general policy is equivalent to finding a chosen-alternative set which maximizes social welfare, subject to a disutility constraint.

Proposition 5.4. *Let \mathbf{z} be an optimal general policy with disutility threshold Q . Any other policy \mathbf{z}' inducing the same alternatives is also an optimal general policy, independent of the actual value of the single incentives or taxes proposed.*

The following proposition shows that the personalized-incentive policy, considered previously, is still relevant in this more general framework. The proposition states that, for any disutility threshold Q , it is possible to find an optimal policy which is a personalized-incentive policy.

Proposition 5.5. *For any $Q \geq 0$, any optimal personalized-incentive policy \mathbf{y} with budget Q (Definition 3.9) is also an optimal general policy with disutility threshold Q (Definition 5.3).*

The following corollary states that the social welfare bound for personalized-incentive policies (Theorem 4.2) is equivalent for general policies.

Corollary 5.6. *Let us run Algorithm 1 with budget Q to construct a personalized-incentive policy. The social welfare $B(Q)$ we obtain is boundedly close to the optimum $B^*(Q)$, obtainable with an optimal general policy with disutility threshold Q :*

$$B^*(Q) - B(Q) \leq \tilde{\epsilon}_{s,t} \cdot (Q - \tilde{Q}).$$

5.2. Optimal Enforcement and Proportional Tax-Subsidy Policy

In this section, we show how Algorithm 1 can be used, in conjunction with Propositions 5.5 and Corollary 5.6, to compute an enforcement policy and a proportional tax-subsidy policy boundedly close to the optimum.

We provide a numerical comparison between these policies in Section 7.6.

5.2.1. Enforcement Policy

With enforcement policies, the regulator constrains the individuals to choose an alternative among a subset of their choice set. In the most extreme case, the individuals can choose only one alternative.

Let \mathbf{y} be the personalized-incentive policy returned by Algorithm 1, for a budget Q . Consider now a policy \mathbf{z} enforcing the individual to choose the same alternative that they would choose under the policy \mathbf{y} , i.e., $z_{i,j} = 0$ if $j = j_i^*(\mathbf{y})$ and $z_{i,j} = -\infty$ otherwise, for any individual i and any alternative j .

Proposition 5.7. *The enforcement policy \mathbf{z} constructed above is boundedly close to an optimal general policy with disutility constraint Q .*

The bound is the same as Corollary 5.6.

5.2.2. Proportional Tax-Subsidy Policy

We consider here policies \mathbf{z} for which the monetary transfers are proportional to the social indicator of the alternatives, that is

$$z_{i,j} = \tau \cdot (b_{i,j} - A_i), \quad \forall i, j, \quad (21)$$

where $\tau > 0$ is the *tax-subsidy level* and $A_i \in \mathbb{R}$ is an individual-specific *baseline social-indicator*, set by the regulator. We call these policies *proportional tax-subsidy policies*. Observe that, for any individual i , her alternatives $j \in \mathcal{N}_i$ such that the social indicator is below the baseline are taxed (i.e., $b_{i,j} < A_i \Rightarrow z_{i,j} < 0$). Conversely, alternatives $j \in \mathcal{N}_i$ having social indicator above the baseline are subsidized (i.e., $b_{i,j} > A_i \Rightarrow z_{i,j} > 0$). The baseline social-indicators A_i can vary from individual to individual. However, we impose that the tax-subsidy level τ is the same for everyone. In this sense, we consider that these policies are not personalized.

Observe from equations (1) and (21) that, considering any individual i , if we vary the baseline A_i the variation of the utility $U_{i,j}$ is the same for all alternatives $j \in \mathcal{N}_i$. Hence, the value of A_i does not impact the choice of i . It simply represents a monetary transfer between the individual and the regulator. More precisely, setting low A_i favours transfers from the regulator to individuals, thus increasing individual utilities, to the detriment of the regulator. On the other hand, setting high A_i , favours the revenue of the regulator, to the detriment of the utility of the individuals.

Note that, if $b_{i,j}$ represents a negative externality (as in Section 7), then the tax-subsidy policy defined above is equivalent to a Pigouvian tax if τ is set to be equal to the external marginal cost of the externalities and $A_i = 0$. For instance, it has been estimated (Quinet et al. 2009), that the social cost of 1 ton of CO₂ is 100 euros. Then, if $b_{i,j}$ represents CO₂ emissions (in tons), the Pigouvian tax would be a proportional tax-subsidy policy with $\tau = 100$ euros.

The following theorem shows that we can use Algorithm 1 to compute a proportional tax-subsidy policy that is boundedly close to the theoretical optimum.

Theorem 5.8. *We can construct a proportional tax-subsidy \mathbf{z} under a certain disutility threshold Q as follows. Run Algorithm 1 with budget constraint Q and let $\tilde{e}_{s,t}$ be the incremental efficiency of the split item given as output.*

The proportional tax-subsidy policy \mathbf{z} , defined as in equation (21), with tax-subsidy level

$$\tau = 1/\tilde{e}_{s,t} \quad (22)$$

achieves a social welfare that is boundedly close to the optimal general policy with disutility threshold Q . The bound is the same as Corollary 5.6.

Up to this point, we have shown that Algorithm 1 is capable of computing nearly optimal personalized-incentive, enforcement, and proportional tax-subsidy policies, as defined by optimization problem (20). To find a policy that maximizes social welfare under budget constraints, exclusively involving positive transfers (without taxation), as presented in optimization problem (7), we have established that Algorithm 1 generates a nearly optimal personalized-incentive policy (Theorem 4.2). The comparison of the personalized-incentive policy and the proportional-subsidy policy within this framework is further discussed in Appendix B.

6. Imperfect Information

Assumption 3.4 states that the regulator knows perfectly the utility of the individuals. With the increasing availability of individual-level data (see e.g., Araldo et al. 2019, Zhu et al. 2020), this assumption is deemed plausible. For example, Zhu et al. (2020) propose a mechanism design to elicit individual preferences from additional data collected. As stated in the introduction, our algorithm is also adapted to different contexts, where the perfect-information assumption might seem more realistic, such as firms tracking users and proposing them personalized discounts. Nevertheless, we show in this section that our algorithm is still relevant when the utility is imperfectly known.

From discrete-choice theory (Anderson, de Palma, and Thisse 1992), we assume that intrinsic utility of alternative j of individual i is composed of a deterministic part $\hat{V}_{i,j}$ and a random part $\epsilon_{i,j}$: $V_{i,j} = \hat{V}_{i,j} + \epsilon_{i,j}$. We assume that the regulator knows the deterministic part $\hat{V}_{i,j}$ of the utility but not the random part $\epsilon_{i,j}$.

Under this assumption, the regulator cannot compute the minimum incentive amount needed to induce individual i to shift from her default alternative $j_i^*(0)$ to another alternative j , using directly equation (6). A heuristic solution would be to set the incentive amount equal to the expectation of the utility difference between the two alternatives, given that $j_i^*(0)$ is the default alternative chosen when there is no incentive. In this case, the incentives $\{y_{i,j}\}_{j \in \mathcal{N}_i}$ proposed by the regulator to individual i , to convince her to shift to alternative j , are such that $y_{i,j'} = 0$, for any $j' \neq j$, and

$$y_{i,j} = \mathbb{E}(V_{i,j_i^*(0)} - V_{i,j} | V_{i,j_i^*(0)} > V_{i,j}) = \hat{y}_{i,j} + \mathbb{E}(\epsilon_{i,j_i^*(0)} - \epsilon_{i,j} | \epsilon_{i,j_i^*(0)} - \epsilon_{i,j} > -\hat{y}_{i,j}), \quad (23)$$

where $\hat{y}_{i,j} = \hat{V}_{i,j_i^*(0)} - \hat{V}_{i,j}$ is the difference in the deterministic part of the utility, known to the regulator.

Given an individual i and an alternative $j \in \mathcal{N}_i$, if the regulator proposes the incentive $y_{i,j}$, as defined by equation (23), then individual i has a positive probability to refuse the incentive. Hence, the expenses of the regulator may be smaller than the total incentive amount proposed.

Algorithm 1 can be used to compute a personalized-incentive policy under imperfect information, by defining new weights

$$w_{i,j} = \mathbb{E}(V_{i,j_i^*(0)} - V_{i,j} | V_{i,j_i^*(0)} > V_{i,j}).$$

At each iteration of the algorithm, the regulator proposes the incentive $w_{i',j'}$ to individual i' for alternative j' , where $[i',j']$ is the pair of individual and alternative selected by the algorithm. The regulator observes the response of the individual to the incentive. If the individual accepts the incentive, it decreases the budget by the

incentive amount. The regulator keeps proposing incentives one by one until his budget is depleted.

Note that, if an individual i accepts an incentive $y_{i,j}$ for alternative $j \in \mathcal{N}_i$, the regulator can still propose her, later, an incentive $y_{i,j'}$ for another alternative $j' \in \mathcal{N}_i$. If the individual refuses the second incentive $y_{i,j'}$, she still receives the first incentive $y_{i,j}$.

In Section 7.7, we apply the policy presented above to our case study and compare it to the case with perfect information, assuming that random terms are Gumbel's distributed. The following proposition gives the exact expression of the incentives (23), in case of Gumbel's distributed random terms.

Proposition 6.1. *Assume that the random terms are i.i.d. Gumbel's distributed with scale parameter μ . Then, the incentive amount from equation (23) is*

$$y_{i,j} = \mu \frac{1 + e^{\hat{y}_{i,j}/\mu}}{e^{\hat{y}_{i,j}/\mu}} \ln \left(1 + e^{\hat{y}_{i,j}/\mu} \right) \geq 0.$$

7. Numerical Results in an Application to Mode Choice

In this section, we simulate an application of our personalized incentive policy to a scenario related to mode choice of individuals commuting to their workplace. We consider a regulator willing to employ a limited monetary budget in order to promote eco-friendly modes of transportation. The goal of the regulator is to reduce CO₂ emissions. We compute the reduction in CO₂ emissions achieved via the personalized-incentive policy of Algorithm 1 and compare it with enforcement, proportional taxation and non-personalized-incentive policies.

Our approach is as follows. After describing the census data used to build the simulation scenario (Section 7.1), we estimate a Multinomial Logit model for mode choice (Section 7.2). Then, using the previous estimates, we simulate the utility of a home-work trip for a group of individuals, for all the modes of transportation considered (Section 7.3). We then approximate the CO₂ emissions for these same trips (Section 7.4) and approximate the optimal personalized-incentive policy using Algorithm 1 (Section 7.5). We then study the modal shifts induced by such policy and the gain in CO₂ emissions achieved. We conclude the numerical results by comparing our personalized incentive policy with other policies (enforcement, taxation, flat incentives – Section 7.6) and by evaluating its performance in case of imperfect information (Section 7.7).

7.1. Data

We use census data from the French statistics institute INSEE, regarding households surveyed between 2015 and 2019. We restrict the dataset to households whose home and workplace are in the Rhône department, which includes Lyon and its suburbs (about 222 000 households in total). Observed variables include city- or district-level home and work location, main mode of transportation used for commuting, and some socio-demographic variables. The modes of transportation are divided in five categories: car, public transit, walking, cycling and motorcycle. Appendix C provides a detailed description of the data.

7.2. *Multinomial Logit Model*

Using the census data, we estimate a Multinomial Logit model for the mode choice of the individuals. We consider five exogenous variables specific to the individual (age, gender, number of cars owned per employee in the household and professional occupation) and one exogenous variable which is specific to both the mode of transportation and the individual (travel time). The number of cars owned is supposed to only impact the utility of commuting by car. Following Inoa, Picard, and de Palma (2015), we also include interaction variables between travel time and socio-demographic variables. Details on how travel time is computed are provided on Appendix D. To estimate the utility of the round trip to work, travel times are doubled (we assume that the modes of transportation for the trip back and forth are the same).

Note that public transit is excluded from the choice set of the individuals whose commute to work cannot be performed by public transit ($\simeq 16\,000$ individuals, see Appendix D for more details).

The four other modes of transportation (car, walking, cycling and motorcycle) are assumed to be in the choice set of all the individuals. This is a strong assumption. A regulator willing to deploy the personalized-incentive policy in practice could improve the precision of the model by using individual-specific data for vehicle ownership in order to remove some modes of transportation from the choice set of an individual, if she does not own the corresponding vehicle.

Table 2 provides the results of the Multinomial Logit model, estimated with the R package `mlogit`. The most frequent categories are used as reference category (car for the mode of transportation, man for the gender and employee for the occupation).

The results are consistent with the literature on commute mode choice. For example, we find that being young and male increases the probability to commute by cycling, consistently with the literature review of cycling mode choice from Muñoz, Monzon, and Daziano (2016). We also find that the coefficient of travel time is larger for public transit than for car which suggests that the value of time for commutes by public transit is slightly smaller than for commutes by car. This is coherent with the meta-analysis on the value of travel-time in France from Wardman et al. (2012).

7.3. *Simulating Utilities*

We consider a regulator whose goal is to reduce the CO₂ emissions due to commute trips, by distributing incentives to the population described in the data (about 222 000 individuals). To apply Algorithm 1, the regulator needs to know the utility and the CO₂ emissions of each individual, for each mode of transportation. We describe in this subsection how we estimate them. Note that our estimations are individual specific.

In the context of the mode choice application, Assumption 3.2 implies that the utility of an individual when commuting by car or public transit does not depend on how many other individuals commute by car or by public transit. However, individuals switching from car to another mode of transportation might reduce road congestion. The congestion reduction, if perceived by the individuals, can in turn affect their utilities and thus modify the optimal incentive amounts. In an iterative procedure where the policy is applied day after day, the optimal policy can be recomputed each time the expected congestion changes. In this case, Assumption 3.2 can be replaced by the following assumption: When the incentives for day t are proposed to the individuals, they do not anticipate the impact that the policy will have on congestion for day t . This assumption is more realistic as it allows the utility of the individuals to depend

Table 2. Multinomial Logit model of mode choice

	(1) car	(2) public_transit	(3) walking	(4) cycling	(5) motorcycle
constant		2.7709*** (0.0395)	2.8659*** (0.0488)	1.1340*** (0.0509)	-0.7284*** (0.0773)
age		-0.0150*** (0.0008)	-0.0026*** (0.0009)	-0.0139*** (0.0010)	-0.0019 (0.0015)
woman		0.5349*** (0.0194)	0.4361*** (0.0248)	-0.3882*** (0.0242)	-1.6909*** (0.0527)
car_per_indiv	1.2138*** (0.0161)				
car_per_indiv>0	1.5604*** (0.0245)				
occupation: farmer		-3.9054*** (0.4140)	-1.0434*** (0.2012)	-2.3653*** (0.5073)	-0.8798** (0.4400)
occupation: artisan		-1.7023*** (0.0525)	-1.2153*** (0.0566)	-0.7848*** (0.0651)	-0.2261*** (0.0841)
occupation: executive		0.1522*** (0.0255)	0.2031*** (0.0327)	1.1710*** (0.0337)	0.2986*** (0.0575)
occupation: intermediate		-0.2283*** (0.0242)	-0.1447*** (0.0311)	0.4259*** (0.0349)	-0.0060 (0.0584)
occupation: blue-collar		-0.7579*** (0.0318)	-0.9691*** (0.0413)	-0.4808*** (0.0467)	-0.0259 (0.0616)
travel_time	-1.6281*** (0.0530)	-1.1746*** (0.0480)	-2.1032*** (0.0492)	-2.8474*** (0.0581)	-3.2075*** (0.0968)
travel_time × age			-0.0026** (0.0010)		
travel_time × woman			-0.1134*** (0.0266)		
travel_time × occupation: farmer			1.1027*** (0.3621)		
travel_time × occupation: artisan			-0.0763 (0.0918)		
travel_time × occupation: executive			-0.3671*** (0.0354)		
travel_time × occupation: intermediate			-0.1986*** (0.0330)		
travel_time × occupation: blue-collar			0.2623*** (0.0403)		

Reference category is male employee
Travel time is expressed in hours
Standard errors are reported in parentheses
*** p<0.01, ** p<0.05, * p<0.1

on the congestion they observed during the previous days. The results presented below are the results for the first day the policy is applied. Results for the following days could be computed by interacting iteratively the algorithm with a transport simulator to update the congestion anticipated on the roads.

Following the Multinomial Logit theory, we assume that utility of alternative j of individual i is composed of a deterministic part $\hat{V}_{i,j}$ and a random part $\epsilon_{i,j}$: $V_{i,j} = \hat{V}_{i,j} + \epsilon_{i,j}$. The deterministic part of the utility can be computed using the estimates from Table 2. As for the random part, we simulate random draws from a random variable with standard Gumbel distribution (see Appendix E). In accordance with Assumption 3.4, the regulator is assumed to know perfectly both the estimates and the draws and thus the utilities. We relax this assumption in Section 7.7, where we provide results where the random draws are unknown to the regulator.

To normalize the utility in monetary units, we compare the value of travel time by car from our regression (expressed in utility units) with the value of travel time by car in France from the literature (expressed in euros). We compute the value of travel time by taking the opposite of the average marginal effect on utility of increasing the travel time of the individuals by one hour.³ We find an average value of time of 1.88 utility units per hour.

Previous studies (Wardman et al. 2012) have shown that the value of travel time, for car commuters, in France, is about 9.17 euros per hour. This would imply that, in our estimates, one utility unit corresponds to $\mu = 9.17/1.88 = 4.88$ euros. In the following, we assume that the utility is normalized in monetary units, i.e., the values of $\hat{V}_{i,j}$ are multiplied by μ . Note that it implies that the random variables $\epsilon_{i,j}$ follows a Gumbel distribution with scale parameter μ .

7.4. Computing the Social Indicator

The regulator wants to reduce greenhouse gas emissions. The social indicator associated to the mode of transportation j of individual i is the reduction in CO₂ equivalent of greenhouse gas emissions generated during the trip of i performed with mode j , with respect to the emissions of the default mode. To compute CO₂ emissions for each individual and each mode of transportation, we take the distance of the fastest path between the individual’s home and workplace and we multiply this distance with the CO₂ emissions equivalent per kilometre for the mode of transportation, using open-sourced data from the French agency ADEME (Agence de l’Environnement et de la Maîtrise de l’Énergie).⁴

For car, we use the CO₂ emissions of a passenger car with average motorization (0.193 kilogram of CO₂ per kilometre). That is, we assume that the CO₂ emissions per kilometre are the same for everyone. We pinpoint that this assumption may lead to some imprecision in the calculation of the actual CO₂ reduction. The application could be improved by using detailed data on the characteristics of the vehicle used by each individual.

As Chester, Horvath, and Madanat (2010), we adopt a disaggregated view of CO₂ emissions from public transit. We consider that the overall CO₂ generated by transit

³For example, referring to the coefficients of Table 2, for a 40-year old male employee, the value of travel time by car is

$$-\left(\underbrace{-1.6281}_{\text{coef. of travel_time (car)}} + \underbrace{-0.0026}_{\text{coef. of travel_time} \times \text{age}} \cdot \underbrace{40}_{\text{age}} \right) = 1.7321 \text{ utility units/hour.}$$

⁴<https://www.bilans-ges.ademe.fr/en/>

Table 3. Our calculations of CO₂ emissions in the Rhône department.

Daily CO ₂ emissions (all home-work and work-home trips)	595.26 tons of CO ₂
Total CO ₂ emissions in one year (200 working days)	119 050 tons of CO ₂
Average yearly individual CO ₂ emissions	0.54 tons of CO ₂

vehicles is shared among all travellers making trips within transit, proportionally to the kilometres travelled. In other words, each trip on transit produces a quantity of CO₂ emissions equal to the number of kilometres travelled multiplied by the average CO₂ emissions per kilometre per passenger, assuming average and constant occupancy rate. The average CO₂ emissions per kilometre per passenger vary according to the mode of transportation used (e.g., bus, tramway or metro). The mode of transportation taken for the fastest path are used to compute CO₂ emissions. For multi-modal public-transit trips (e.g., bus then tramway), the CO₂ emissions are computed according to the distance travelled by each mode of transportation.

CO₂ emissions for walking and cycling trips are set to zero. Hence, for each individual, the two alternatives corresponding to walking and cycling differ only in the intrinsic utility.

Under the previous assumptions, we calculate the CO₂ emissions reported in Table 3, which results in 0.54 ton of CO₂ yearly per individual in the Rhône department. This number is close to the publicly known estimation for the entire France: in 2007, the average French worker emitted 0.64 ton per year because of his/her home-work trips (Levy and Le Jeannic 2011).

Assumption 3.3 implies that the CO₂ emissions due to the commuting trip of an individual must be independent from the mode of transportation chosen by the other commuters. This assumption is violated because changes in congestion induce a slight change in CO₂ emissions emitted. Assumption 3.4 implies that the regulator knows perfectly the CO₂ emissions of the trips. This assumption is somewhat questionable since we only approximate CO₂ emissions. More accurate estimates of CO₂ emissions are potentially available using a detailed emission model coupled with a transport simulator (see e.g., Rahman et al. 2021), but we wished to keep this application simple. Nevertheless, measurement errors for CO₂ emissions are not as worrying as measurement errors for utility as we can assume that, if such errors are unbiased, they cancel out. We will observe in Section 7.7 that the errors are much more severe when utilities are imperfectly known, as some individuals might reject the incentives, which leads to a suboptimal allocation.

7.5. Calculation of the Personalized-Incentive Policy

We consider a large-scale scenario with more than 200 thousands individuals and over 1 million alternatives (Appendix C). We consider a policy in which the regulator proposes, each day, incentives to the individuals before their home-work trip. The incentives are given conditional on the mode of transportation chosen for the round trip to work, thus the social indicator of an alternative is the reduction in CO₂ emissions for the trip *back and forth*, with respect to the default alternative. The budget represents the daily amount available to the regulator for incentives.

First, we run Algorithm 1 with a daily budget of 3000 euros and we plot the maximum social welfare curve (see Figure 2). The maximum social welfare curve is an




Figure 2. Maximum social welfare curve, up to a daily budget of 3000 euros. **Note:** The social welfare corresponds to the reduction in CO₂ emissions due to the personalized-incentive policy.

increasing step function (steps are small and thus not visible). Consistently with Proposition 4.7, the slope of each step is non-increasing, which gives the curve a concave curvature.

Quinet et al. (2009) predict that the carbon price in France would be of 100 euros per ton of CO₂ in 2030. It is thus reasonable to assume the regulator is interested in finding an incentive policy such that, for every 100 euros spent in incentives, pollution is reduced by at least a ton of CO₂. To this aim, the regulator can observe the curves of Figure 3, which plots the inverse of the incremental and overall efficiency (the $\tilde{e}^{[k]}$ and $e^{[k]}$ of Definition 4.6), with respect to the budget $Y^{[k]}$ allocated by the algorithm at each iteration k . Thanks to Proposition 4.7, $1/e^{[k]}$ and $1/\tilde{e}^{[k]}$ increase with $Y^{[k]}$, as we proceed with the iterations of the algorithm. Thanks to this monotonicity, the regulator can apply one of the following two criteria to fix the budget to invest. It could run Algorithm 4 and stop it when $1/e^{[k]}$ equals 100 euros per ton of CO₂. Alternatively, it can stop the Algorithm when $1/\tilde{e}^{[k]}$ equals 100 euros per ton of CO₂. From Figure 3, we observe that with the first criterion the regulator would need to invest about 1800 euros per day, and about 500 euros with the second criterion. In our opinion, both criteria would make sense, and the preference over one of them is a political choice.

We now set the budget of the regulator to $Q = 1800$ euros. Running Algorithm 1 with this budget required about 3500 iterations and took about 6 seconds (with Python, on a computer with an Intel i5-8350U 1.7GHz and 24GB of memory). The algorithm allocates practically all the budget (1798.59 euros). We find that 1.57% of individuals received incentives and changed transportation mode, which results in a reduction of CO₂ emission by 18 tons of CO₂ per day (3.00 % of total CO₂ emissions). Thus, this policy would cost on average 100.61 euros for each ton of CO₂ prevented.

Despite the small incentives, the reduction in CO₂ emissions is considerable. Indeed, among the individuals who received incentives, the average amount of incentives is 0.52 euros per individual, for an average daily reduction in CO₂ emissions of 5 kilograms. Given that the average travel time for the round trip to work by car is




Figure 3. Cost of the policy in euros per ton of CO₂ prevented as a function of the daily budget.

0.67 hour (40 minutes) and the value of time by car is 9.17 euros per hour, an incentive of 0.52 euros represents 8.5% of the value of time for the average car trip ($0.52/(9.17 \times 0.67) = 0.085$). Recall that alternatives providing a large reduction in CO₂, while requiring small incentive, have a high efficiency. Hence, the algorithm selects first shifts achievable with a small incentive, i.e., where the individual is almost indifferent between the two alternatives, which however have a large difference in CO₂. Figure 4 shows the distribution of the incentive amount and the CO₂ reduction for the incentivized individuals. For most incentives, the amount proposed to individuals is below 1 euro (incentives with a larger amount are not efficient enough, unless the CO₂ reduction is very high).

Figure 5 compares mode share before and after the policy. Most individuals who received incentives are individuals who commuted by car and were induced to commute by public transit (1.2% of all individuals, 74% of individuals who received incentives). The share of individuals commuting by car decreased by 2.4%, while public transit ridership increased by 4%.

We now compute a bound of the optimality gap, i.e., the maximum additional CO₂ savings we would achieve if we could use a theoretical optimal policy instead of resorting to Algorithm 1. To do so, we apply Theorem 4.2. Since the incremental efficiency of the split item returned by the algorithm is $\tilde{e}_{s,t} \simeq 5$ kilograms of CO₂ per euro and the unused budget is $Q - \tilde{Q} = 1.41$ euros, an optimal policy would reduce just $5 \cdot 1.41 \simeq 7$ kilograms more than Algorithm 1, which is negligible compared to the total CO₂ emissions reduction of 18 tons provided overall.

7.6. Comparison with Other Policies

In Section 7.5, we evaluated the performance of the personalized incentive policy calculated by Algorithm 1. If the regulator is not limited by his budget or if he is allowed to use taxation or enforcement, then other policies are potentially available (if they




Figure 4. Distribution of incentive amount and CO₂ reduction for the incentives given in one day with budget $Q = 1800$ euros. The slope of the black line represents the incremental efficiency of the split item returned by the algorithm, $\bar{e}_{s,t} = 5$ tons of CO₂ / euro. Note that all points are above the line because their incremental efficiency is larger. The histogram above represents the distribution of the incentive amounts. The histogram on the right represents the distribution of the CO₂ reduction for the incentives.




Figure 5. Evolution of mode share before and after the policy. 1.163 % of individuals were given incentives to shift from car to public transit, 27.29 % of individuals commuted by public transit before the policy and were not induced to shift.

Table 4. Summary of policies.

Policy \mathbf{z}	Expenses (euros) $Y(\mathbf{z})$	Ind. utility (euros) $\Delta U(\mathbf{z})$	Disutility (euros) $\delta(\mathbf{z})$	CO ₂ reduction (tons) $B(\mathbf{z})$
Personalized incentives	1798.59	0	1798.59	17.878
Enforcement	0	-1798.59	1798.59	17.878
Proportional tax	-114 368.89	-116 167.48	1798.59	17.878
Proportional subsidies	3596.97	1798.38	1798.59	17.878

are politically accepted), as shown in Section 5. In this section, we compare the results of the personalized incentive policy with three other policies: an enforcement policy, a proportional taxation system and a proportional subsidy system.

Aggregate results for these policies are provided in Table 4. The policies are defined so that they induce the same choices for the individuals, using results from Section 5.2. Therefore, they provide the same reduction in CO₂ emissions and, from Proposition 5.2, they have the same disutility. However, they differ in their cost for the regulator and the variation in individual utilities implied. It should be noted that the best policy to implement depends on social, political or juridical constraints.

In these results, we fix the disutility threshold to 1798.59 euros, which corresponds to the incentive \tilde{Q} actually spent by Algorithm 1 when we set the budget to $Q = 1800$ euros. This means that if we run the algorithm setting a budget of 1798.59 euros, it spends it all and, thanks to Corollary 5.6, the resulting policy is optimal under a budget constraint of 1798.59 euros.

7.6.1. Enforcement Policy

Thanks to Proposition 5.7, the regulator can compute an enforcement policy \mathbf{z} that is optimal for a disutility threshold of 1798.59 euros, by simply ‘imitating’ the personalized incentive policy \mathbf{y} , i.e., by inducing the same alternatives as \mathbf{y} . In order to do so, the regulator bans all the other alternatives, i.e., any alternative j such that $j \neq j_i^*(\mathbf{y})$ is banned. Obviously, it is not necessary to ban any alternative j if $j_i^*(\mathbf{y})$ is preferred to j , in absence of policy, i.e., $V_{i,j} < V_{i,j_i^*(\mathbf{y})}$. Therefore, only the individuals receiving incentives under policy \mathbf{y} suffer bans with \mathbf{z} , which correspond to only 1.57% of the population.

Contrarily to the personalized-incentive policy, the enforcement policy does not cost any money to the regulator (apart from eventual transaction costs) but it decreases individual utilities by 1798.59 euros. Moreover, the 1.57% of individuals impacted by the ban may perceive that they are inequitably penalized with respect to the others. Hence, the enforcement policy might be less accepted by the population. Still, this policy is well adapted to the context of imperfect information as it ensures that the individuals always choose the alternative wanted by the regulator.

7.6.2. Proportional Tax

The proportional tax policy is computed from equation (21), using the tax level given by equation (22) (Theorem 5.8). For each individual i , the baseline social-indicator is set to $A_i = 0$, so that the transfer amount for alternative j is equal to the opposite of the CO₂ emissions of transportation mode j , and is thus always non-positive.

Since the taxation provides revenues, the regulator is not constrained by his budget anymore. However, taxation negatively impacts the utilities of the individuals, and is thus limited by political constraints, which we model by imposing that the disutility

of the policy must be below a threshold $Q = 1798.59$ euros.

The tax must be paid by all individuals commuting either by car, public transit or motorcycle (about 86 % of individuals), even if the tax does not affect their choice. This is different from the personalized-incentive policy, for which only 1.57 % of individuals are impacted. This explains why the amount of taxes collected is about hundred times larger than the amount of incentive needed to reach the same reduction in CO₂ emissions (see Table 4). A taxation policy is particularly penalizing for inelastic individuals who cannot shift to a more eco-friendly alternative, e.g., because they are living far from their workplace, or in a place with no transit offer. Therefore, a taxation policy is much less acceptable than an incentive one.

On the other hand, the tax policy is not individual-specific, which means that it requires less information (knowledge of individual utilities is required to compute the tax level from (22) but the tax level does not change much under imperfect information and so the policy is still efficient).

7.6.3. *Proportional Subsidy / Tripod*

The proportional-subsidy policy is computed in the same way as the proportional-tax policy but using $A_i = b_{i,j_i^*}(0)$, for any individual i . This means that, (i) if the individual uses the same mode with and without the policy (her default mode), no monetary transfer occurs; (ii) if the individual switches to a mode with less CO₂ emissions, she will receive a subsidy proportional to the reduction in CO₂ emissions compared to the default mode; (iii) if the individual switches to a mode involving more CO₂ emissions, she will get taxed (and thus it is easy to show that individuals never choose to do so).

Observe that the proportional-subsidy policy considered in this section, like our personalized-incentive policy, is personalized as it depends on the alternative chosen by default by the individuals. These two policies could thus be perceived as unfair by individuals who initially choose environmentally-friendly alternatives (like walking or bicycling), as they would not receive any incentive. In contrast, individuals who initially did not choose environmentally-friendly alternatives would receive incentives when switching to walking or bicycling. To balance this, it is worth considering that individuals who already walk or bike might highly appreciate policies aimed at reducing CO₂ emissions. Raising awareness about the policy's benefits and implementing additional measures like developing bike infrastructure could contribute to its acceptance among the population.

Such a policy is very close to the Tripod policy (Araldo et al. 2019), in which individuals are being offered tokens for choosing alternatives with less CO₂ emissions, where the token amount is proportional to the reduction in CO₂ emissions. The two main differences between the proportional-subsidy policy presented here and Tripod are that (i) we assume that the set of individuals traveling is known in advance, while Tripod is designed for dynamic scenarios where individuals are showing up stochastically; (ii) the subsidy level τ is set optimally (using Theorem 5.8), while in Tripod it is computed using simulation-based prediction.

Both the personalized-incentive policy and the proportional-subsidy policy are adapted to cases where the regulator is endowed with a limited budget that he must use as efficiently as possible to reach some objective. In such cases, our personalized-incentive policy performs better than the proportional-subsidy policy: the personalized-incentive policy can reach the same level of CO₂ emissions reduction than the proportional-subsidy policy with around half the budget (see Table 4). The reason is that the incentive amounts are based on the willingness-to-accept of




Figure 6. CO₂ emissions reduction and average incentive amount under imperfect information.

the individuals, thus getting the most additional social benefit out of every additional unit of incentive spent. The proportional-subsidy policy is instead limited to an incentive amount depending only on the CO₂ emissions of the alternative. However, the greater effectiveness of the personalized-incentive policy is due to the compromise of (i) requiring more information about individual preferences and (ii) potentially facing discontent among the population due to differential treatment of individuals.

7.7. Imperfect Information

We show in this section the performance of our allocation policy when the regulator has imperfect information about individual utilities. In this case, the allocation policy is computed as in Section 6. Using the values of the random variables $\epsilon_{i,j}$ drawn previously, we can check whether individuals accept the incentives proposed to them. If the incentive is accepted, the expenses of the regulator increase by the incentive amount. The policy stops when the daily budget of 1800 euros is depleted.

Figure 6 shows the CO₂ emissions reduction and the average incentive amount of our personalized-incentive policy under imperfect information, for different values of μ (the scale parameter of the random terms). Figure 7 shows the number of incentives proposed and accepted for the same values of μ . Note that μ can be interpreted as the level of information available to the regulator. For low values of μ , the random terms are small and thus the incentives proposed are close to the incentive amount proposed under perfect information (i.e., the minimum amount required to induce the individuals to switch mode). The incentive amounts are small and many incentives are proposed before the budget is depleted, leading to a large reduction in CO₂ emissions. For large values of μ , the random terms are large and thus the incentive amount is often disconnected from the minimum amount each individual would accept, which implies that many incentives will be either refused or will be accepted but with a large amount compared to the minimum required. Therefore, the budget gets depleted with




Figure 7. Number of incentives proposed and accepted and acceptance rate under imperfect information.

fewer incentives given and thus a smaller reduction in CO₂ emissions is achieved.

Interestingly, Figure 7 shows that the share of incentives accepted is practically constant with μ (except for the smallest value $\mu = 10^{-8}$) and close to 60%. These acceptance ratios can be compared with the ratio of 68% obtained by Zhu et al. (2020) with an online experiment.

These results could be improved by learning from the responses of individual i to the incentives proposed earlier in order to compute the incentives that will be proposed to her for other alternatives, in a similar way to Zhu et al. (2020). For example, if the regulator observes that individual i refused the incentive to shift from car to walking, he learns information on the random term of the utility for car of individual i .

Also, if it is not possible to propose incentives to individual i for different alternatives consecutively, the regulator could propose incentives for multiple alternatives simultaneously. These extensions cannot be carried out with Algorithm 1. Additional work is needed to improve the personalized-incentive policy under imperfect information.

8. Conclusion

This paper explores a new system of personalized incentives. The agents face a discrete set of alternatives, and make independent discrete choices. We consider situations where an individual utility for an alternative does not coincide with the social utility of this alternative. Such situations call for regulation or State intervention. The idea is to determine the optimal incentives to be provided to each individual to alter their choices in order to better align individual benefits and the Principal benefits (note that the Principal can be any regulator). The regulator is assumed to have a fixed budget for the incentives. Even if individuals make independent choices, the computation of the incentives to be provided has to consider all individuals' preferences, so the problem is combinatorial. We provided in this paper an algorithm to optimally distribute

individual incentives given a budget constraint in order to maximize the social utility or the social welfare function.

In 2021, this incentive system may be somewhat in advance. Nowadays, individual information is gathered via GPS, social networks and the Internet of things. This is precious information, which can potentially be used to optimally compute the optimal set of incentives, and thus to better manage Society. (Privacy issues are ignored here, which does not mean they are not important.)

Besides, humans remain unpredictable. There is still (and hopefully for some time) some margin of freedom as far as to what people decide. The recent pandemic shows that individuals or governments remain unpredictable (Zhang 2020) and that the right set of incentives remains hard to determine. As a consequence, individual choices are described by the modeller as being probabilistic. Incentives thus change choices up to some probability distribution. While we have just tackled imperfect information in the empirical application, the treatment of imperfect information appears to be tractable. Preliminary computations, with the Logit, the workhorse of discrete choice models, suggest that such an extension is promising, including analytically. Contrarily to the full information case, mainly envisaged in this paper, some incentives may be too large for some individuals (who could select the same choice with a smaller incentive), and this incentive is then inefficient; other incentives may be too small to modify individual choice as expected, and in such a case the incentive is ineffective. The optimal solution makes a compromise between these two sources of imperfection.

In the empirical application, we have ignored congestion. In our defence, let's recall that few commuters receive an incentive, which is a quality of our method. In practice, congestion means that the utility of some individuals can change as other individuals are shifting, which renders the incentive amounts computed *ex-ante* imprecise. We have not solved the current problem with congestion because it is likely to be difficult. But it is not impossible. An iterative procedure alternating the incentive algorithm and the computation of the current level of congestion is promising. Congestion can be treated as a static or dynamic (time of the day dependent) process. Much work remains to be done along this line.

Finally, we have considered so far static choice, i.e., at a given point in historical time. If we consider mode choice, it may be the case that incentive for public transport, for example, will have on the long run an impact on automobile ownership. Moreover, in the medium run, a car left at home can be used by other family members for short trip. Without any intervention, the trend could yield more trips and vehicle cold starts particularly on local roads, especially in places where vehicles continue to rely on internal combustion engines. These examples show the need to also consider the medium and long-run impacts of incentives, by appending a predictive model to the incentive algorithm. There are plenty of roads left to run.

Acknowledgments

We are grateful to Moshe Ben-Akiva for his constructive and valuable feedback, which helped us to more deeply understand the limits and the implications of the personalized incentive policy proposed in this paper, in particular in a dynamic context and in terms of equity. His comments allowed us to better present and frame our approach within the state-of-the-art. We thank Matthieu de Palma, Nathalie Picard and Ravi Seshadri for their helpful comments at various stages of this research. Finally, we thank three anonymous referees for their valuable and constructive comments.

Disclosure statement

The authors report there are no competing interests to declare.

Funding

This work was supported by NSFC-JPI UE under grant “MAAT” (project no. 18356856); ANR under grant ANR-20-CE22-0014.

References

- Anderson, Simon P, André de Palma, and Jacques François Thisse. 1992. *Discrete choice theory of product differentiation*. MIT press.
- Araldo, A., A. Di Stefano, and A. Di Stefano. 2020. “Resource Allocation for Edge Computing with Multiple Tenant Configurations.” In *ACM/SIGAPP Symposium On Applied Computing*, .
- Araldo, A., S.I Gao, R. Seshadri, C. Lima Azevedo, H. Ghafourian, Y. Sui, S. Ayaz, D. Sukhin, and M. Ben-Akiva. 2019. “System-level optimization of multi-modal transportation networks for energy efficiency using personalized incentives.” *Transportation Research Records* 2673 (12).
- Arcep. 2021. “Le plan France Très Haut Débit.” <https://www.arcep.fr/demarches-et-services/collectivites/le-plan-france-tres-haut-debit-pfthd.html>.
- Cao, Nianxia, Swastik Brahma, and Pramod K. Varshney. 2015. “Target tracking via crowdsourcing: A mechanism design approach.” *IEEE Transactions on Signal Processing* 63 (6): 1464–1476.
- Chester, Mikhail V., Arpad Horvath, and Samer Madanat. 2010. “Comparison of life-cycle energy and emissions footprints of passenger transportation in metropolitan regions.” *Atmospheric Environment* 44 (8): 1071–1079. <http://dx.doi.org/10.1016/j.atmosenv.2009.12.012>.
- Clarke, Amanda, and Helen Margetts. 2014. “Governments and citizens getting to know each other? Open, closed, and big data in public management reform.” *Policy & Internet* 6 (4): 393–417.
- Coloni, Alberto, Valentina Ferretti, Alessandro Luè, Alessandra Oppio, Valerio Parusco, and Luca Tomasini. 2017. “Rethinking feasibility analysis for urban development: A multidimensional decision support tool.” In *Computational Science and its Applications - ICCSA*, .
- De Borger, Bruno. 2001. “Discrete choice models and optimal two-part tariffs in the presence of externalities: optimal taxation of cars.” *Regional Science and Urban Economics* 31 (4): 471–504.
- Ettema, Dick, Jasper Knockaert, and Erik Verhoef. 2010. “Using incentives as traffic management tool: empirical results of the peak avoidance experiment.” *Transportation Letters* 2 (1): 39–51.
- Fielder, Andrew, Emmanouil Panaousis, Pasquale Malacaria, Chris Hankin, and Fabrizio Smeraldi. 2016. “Decision support approaches for cyber security investment.” *Decision Support Systems* 86: 13–23. <http://dx.doi.org/10.1016/j.dss.2016.02.012>.
- Ghafelebashi, Ali, Meisam Razaviyayn, and Maged Dessouky. 2023. “Congestion reduction via personalized incentives.” *Transportation Research Part C: Emerging Technologies* 152: 104153.
- Hu, Xianbiao, Yi Chang Chiu, and Lei Zhu. 2015. “Behavior Insights for an Incentive-Based Active Demand Management Platform.” *International Journal of Transportation Science and Technology* 4 (2): 119–133. <http://dx.doi.org/10.1260/2046-0430.4.2.119>.

- Inoa, Ignacio A, Nathalie Picard, and André de Palma. 2015. “Effect of an accessibility measure in a model for choice of residential location, workplace, and type of employment.” *Mathematical Population Studies* 22 (1): 4–36.
- Javaudin, Lucas, Andrea Araldo, and André de Palma. 2022. “Large-scale allocation of personalized incentives.” In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, 4151–4156. IEEE.
- Kellerer, Hans, et al. 2004. *Knapsack Problems*. 1st ed. Springer.
- Kirkpatrick, David G, and Raimund Seidel. 1986. “The Ultimate Planar Convex Hull Algorithm?” *SIAM J. Comput.* 15 (1): 1–18. <http://ecommons.cornell.edu/handle/1813/6417>.
- Knockaert, Jasper, Yin-Yen Tseng, Erik T Verhoef, and Jan Rouwendal. 2012. “The Spitsmijden experiment: A reward to battle congestion.” *Transport Policy* 24: 260–272.
- Levy, David, and Thomas Le Jeannic. 2011. “Un habitant de pôle urbain émet deux fois moins de CO2 que la moyenne pour se rendre à son lieu de travail ou d’études.” *INSEE Première*.
- Li, Tianhao, Peng Chen, and Ye Tian. 2021. “Personalized incentive-based peak avoidance and drivers’ travel time-savings.” *Transport Policy* 100: 68–80.
- Merugu, Deepak, Balaji S Prabhakar, and N Rama. 2009. “An incentive mechanism for decongesting the roads: A pilot program in bangalore.” In *Proc. of ACM NetEcon Workshop*, Citeseer.
- Mohammadivojdan, Roshanak, and Joseph Geunes. 2018. “The newsvendor problem with capacitated suppliers and quantity discounts.” *European Journal of Operational Research* 271 (1): 109–119. <https://doi.org/10.1016/j.ejor.2018.05.015>.
- Muñoz, Begoña, Andres Monzon, and Ricardo A Daziano. 2016. “The increasing role of latent variables in modelling bicycle mode choice.” *Transport Reviews* 36 (6): 737–771.
- Nadarajah, Saralees, and Samuel Kotz. 2005. “A generalized logistic distribution.” *International Journal of Mathematics and Mathematical Sciences* 2005 (19): 3169–3174.
- Quinet, Alain, Luc Baumstark, Joffrey Célestin-Urbain, Hervé Pouliquen, Dominique Auverlot, and Christine Raynard. 2009. “La valeur tutélaire du carbone.” *Rapport du Conseil d’Analyse Stratégique* 16 (5): 9305.
- Rahman, Md Mamunur, Yuan Zhou, Jamie Rogers, Victoria Chen, Melanie Sattler, and Kate Hyun. 2021. “A comparative assessment of CO2 emission between gasoline, electric, and hybrid vehicles: A Well-To-Wheel perspective using agent-based modeling.” *Journal of Cleaner Production* 321: 128931.
- Small, Kenneth A., and Harvey S. Rosen. 1981. “Applied Welfare Economics with Discrete Choice Models.” *Econometrica* 49 (1): 105–130.
- Sun, Jian, Jiyang Wu, Feng Xiao, Ye Tian, and Xiangdong Xu. 2020. “Managing bottleneck congestion with incentives.” *Transportation Research Part B: Methodological* 134: 143–166. <https://doi.org/10.1016/j.trb.2020.01.010>.
- Tang, Yili, Yu Jiang, Hai Yang, and Otto Anker Nielsen. 2020. “Modeling and optimizing a fare incentive strategy to manage queuing and crowding in mass transit systems.” *Transportation Research Part B: Methodological* 138: 247–267.
- Wardman, Mark, Phani Chintakayala, Gerard de Jong, and Diego Ferrer. 2012. “European wide meta-analysis of values of travel time.” *ITS, University of Leeds, Paper prepared for EIB*.
- Xiong, Chenfeng, Mehrdad Shahabi, Jun Zhao, Yafeng Yin, Xuesong Zhou, and Lei Zhang. 2020. “An integrated and personalized traveler information and incentive scheme for energy efficient mobility systems.” *Transportation Research Part C: Emerging Technologies* 113: 57–73.
- Yue, Jia Shuo, Chinmoy V Mandayam, Deepak Merugu, Hossein Karkeh Abadi, and Balaji Prabhakar. 2015. “Reducing road congestion through incentives: a case study.” In *Transportation Research Board 94th Annual Meeting, Washington, DC*.
- Zhang, Junyi. 2020. “Transport policymaking that accounts for COVID-19 and future public health threats: A PASS approach.” *Transport policy* 99: 405–418.

- Zhong, Tao, and Rhonda Young. 2010. “Multiple Choice Knapsack Problem: Example of planning choice in transportation.” *Evaluation and Program Planning* 33 (2): 128–137.
- Zhu, Xi, Feilong Wang, Cynthia Chen, and Derek D Reed. 2020. “Personalized incentives for promoting sustainable travel behaviors.” *Transportation Research Part C: Emerging Technologies* 113: 314–331.
- Zoltners, Andris A., Prabhakant Sinha, and Philip S. C. Chong. 1979. “An Optimal Algorithm for Sales Representative Time Management.” *Management Science* 25 (12): 1197–1207.

Appendix A. Proofs

Proof of Proposition 3.7. Given any policy \mathbf{y} , individual i chooses alternative $j \in \mathcal{N}_i$, with $b_{i,j} > b_{i,j_i^*(0)}$, if

$$V_{i,j} + y_{i,j} \geq V_{i,j'} + y_{i,j'}, \quad \forall j' \in \mathcal{N}_i, \quad (\text{A1})$$

and

$$V_{i,j} + y_{i,j} > V_{i,j'} + y_{i,j'}, \quad \forall j' \in \mathcal{N}_i \setminus \{j\} : b_{i,j} \leq b_{i,j'}. \quad (\text{A2})$$

Indeed, equations (A1) and (A2) ensure that (4) is satisfied.

Let $i \in \mathcal{I}$ and $j \in \mathcal{N}_i$, and consider a personalized-incentive policy \mathbf{y} such that $y_{i,j'} = 0$, for any $j' \neq j$ and $y_{i,j} = V_{i,j_i^*(0)} - V_{i,j}$. Rewriting (A1) and (A2), we can claim that individual i chooses alternative j , if

$$V_{i,j_i^*(0)} \geq V_{i,j'}, \quad \forall j' \in \mathcal{N}_i, \quad (\text{A3})$$

and

$$V_{i,j_i^*(0)} > V_{i,j'}, \quad \forall j' \in \mathcal{N}_i \setminus \{j\} : b_{i,j} \leq b_{i,j'}. \quad (\text{A4})$$

Thanks to equation (5), the personalized-incentive policy \mathbf{y} satisfies equation (A3). It remains to prove that it always satisfies also equation (A4). Suppose by contradiction that there exists an alternative $j' \in \mathcal{N}_i \setminus \{j\}$ such that $b_{i,j} \leq b_{i,j'}$, which does not satisfy equation (A4). Then we would have $V_{i,j_i^*(0)} \leq V_{i,j'}$. By construction $b_{i,j'} \geq b_{i,j} > b_{i,j_i^*(0)}$. This would contradict the definition of default alternative (equation (5)).

At this point of the proof, we have demonstrated the first part of the Proposition, i.e., that, considering an option j such that $b_{i,j} > b_{i,j_i^*(0)}$, a personalized-incentive policy \mathbf{y} such that $y_{i,j'} = 0$, for any $j' \neq j$ and $y_{i,j} = V_{i,j_i^*(0)} - V_{i,j}$ successfully induces individual i to choose alternative j . We now prove the second part of the Proposition.

Observe that, if $y_{i,j} < V_{i,j_i^*(0)} - V_{i,j}$, then $U_{i,j}(\mathbf{y}) = V_{i,j} + y_{i,j} < V_{i,j_i^*(0)}$ and individual i would never prefer j to $j_i^*(0)$. \square

Proof of Theorem 4.2. If we run Algorithm 1 with budget Q , we practically make the same steps as the MCKP-Greedy algorithm (Kellerer et al. 2004, equation (11.8) and Figure 11.2). In line 5 of the aforementioned algorithm, the authors compute an upper bound to the solution of the Multiple Choice Knapsack Problem (9) as

$$ub = \bar{B}^{[k]} + \tilde{b}_{s,t} \cdot (Q - Y^{[k]}) / \tilde{w}_{s,t},$$

where k is the last iteration of the algorithm.

Observing, by the definition of efficiency (13), that $\tilde{e}_{s,t} = \tilde{b}_{s,t}/\tilde{w}_{s,t}$, we get $ub - \bar{B}^{[k]} = \tilde{e}_{s,t} \cdot (Q - Y^{[k]})$. By construction, the theoretical maximum social welfare $B^*(Q)$ of problem (9) is less than or equal to the upper bound ub . Therefore:

$$B^*(Q) - \bar{B}^{[k]} \leq ub - \bar{B}^{[k]} = \tilde{e}_{s,t} \cdot (Q - Y^{[k]}).$$

By construction, $B(Q) = \bar{B}^{[k]}$ and $\tilde{Q} = Y^{[k]}$, which gives the inequality (17) that we want to prove. Such inequality is illustrated in Figures 1 and A1.




Figure A1. Illustration of the upper bound from Theorem 4.2.

□

Proof of Proposition 4.5. To compute the ordered LP-extremes \mathcal{R}_i of the individual i we resort to the method of Kirkpatrick and Seidel (1986) of complexity $O(\sum_{i=1}^m |\mathcal{N}_i| \cdot \log |\mathcal{R}_i|)$. To obtain the set \mathcal{R} , we just need to merge these ordered sets into an aggregated ordered set. This operation has complexity $O(|\mathcal{R}| \cdot \log m)$. The rest of the operations consists in adding to the solution the alternatives in \mathcal{R} , one by one, which has complexity $O(|\mathcal{R}|)$. □

Proof of Proposition 4.7. By construction, Algorithm. 1 gets at each iteration the alternative with the highest incremental efficiency (Line 8). This proves the first part of the claim.

The second part of the claim can be shown geometrically. In Figure A2, we represent the total incentive and social welfare calculated by the algorithm at each iteration.




Figure A2. Incremental efficiency at each iteration.

Observe that the incremental efficiency $\tilde{e}^{[k]}$ is the inclination of the segment connecting $(Y^{[k-1]}, B^{[k-1]})$ to $(Y^{[k]}, B^{[k]})$ and that the efficiency $e^{[k]}$ is the inclination of the segment connecting $(0, 0)$ to $(Y^{[k]}, B^{[k]})$. It becomes then evident that the monotonicity of $\tilde{e}^{[k]}$ implies also the monotonicity of $e^{[k]}$. □

Proof of the Proposition 5.2. From equations (1), (3), (18) and (19), observe that

$$\delta(\mathbf{z}) = \sum_{i=1}^m (V_{i,j_i^*}(0) - V_{i,j_i^*}(\mathbf{z})) \geq 0,$$

□

Proof of Proposition 5.4. Thanks to Proposition 5.2, $\delta(\mathbf{z}') = \delta(\mathbf{z}) \leq Q$. Moreover, the social welfare is also the same, i.e., $B(\mathbf{z}) = B(\mathbf{z}')$, since it only depends on the alternative chosen. □

Proof of Proposition 5.5. Assume, by contradiction, that the optimal personalized-incentive policy \mathbf{y} is not an optimal general policy. This would imply the existence of a policy \mathbf{z} such that $B(\mathbf{z}) > B(\mathbf{y})$ and $\delta(\mathbf{z}) \leq Q$.

Consider now a personalized-incentive policy \mathbf{y}' such that

$$\begin{cases} y'_{i,j} = V_{i,j_i^*(0)} - V_{i,j}, & \text{if } j = j_i^*(\mathbf{z}) \\ y'_{i,j} = 0, & \text{otherwise} \end{cases}.$$

Then, by construction, \mathbf{y}' is such that $j_i^*(\mathbf{y}') = j_i^*(\mathbf{z}), \forall i \in \mathcal{I}$.

Moreover, observe that \mathbf{y}' is such that $B(\mathbf{y}') > B(\mathbf{y})$ and $Y(\mathbf{y}') = \delta(\mathbf{y}') \leq Q$. Therefore, \mathbf{y}' would be a better personalized-incentive policy than \mathbf{y} , which is absurd, since by construction \mathbf{y} is an optimal personalized-incentive policy. □

Proof of Corollary 5.6. Let $B(Q)$ be the social welfare returned by Algorithm 1 for a budget Q . We know from Theorem 4.2 that it is boundedly close to the social welfare $B^*(Q)$ obtained with an optimal personalized-incentive policy, with the following bound

$$B^*(Q) - B(Q) \leq \tilde{e}_{s,t} \cdot (Q - \tilde{Q}).$$

Thanks to Proposition 5.5, $B^*(Q)$ is also the social welfare obtained via an optimal general policy with disutility threshold Q . □

Proof of Proposition 5.7. By construction, the enforcement policy \mathbf{z} induce the same individual alternatives as the personalized-incentive policy \mathbf{y} . Then, thanks to Proposition 5.2, they have the same disutility $\delta(\mathbf{z}) = \delta(\mathbf{y}) \leq Q$ and achieve the same social welfare $B(\mathbf{z}) = B(\mathbf{y})$. Thanks to Corollary 5.6, $B(\mathbf{y})$ is boundedly close to the optimum $B(Q)$, and so is $B(\mathbf{z})$. □

Proof of Theorem 5.8. Let \mathbf{z} be a policy such that $z_{i,j} = \tau(b_{i,j} - A_i)$, with τ given by equation (22) and $A_i \in \mathbb{R}$.

Let \mathbf{y} be the personalized-incentive policy obtained running the algorithm as explained in the statement of this theorem. Thanks to Corollary 5.6, we know that $B(\mathbf{y})$ is such that

$$B^*(Q) - B(\mathbf{y}) \leq \tilde{e}_{s,t} \cdot (Q - \tilde{Q}). \quad (\text{A5})$$

If we prove that $j_i^*(\mathbf{y}) = j_i^*(\mathbf{z}), \forall i \in \mathcal{I}$, we could claim that $B(\mathbf{z}) = B(\mathbf{y})$ and also, thanks to Proposition 5.2, that $\delta(\mathbf{z}) = \delta(\mathbf{y}) \leq Q$. In this case, the bound (A5) would also hold for \mathbf{z} .

To do so, we show that (i) alternative $j_i^*(\mathbf{z})$ is in the set \mathcal{R}_i of the LP-extremes alternatives and (ii) alternative $j_i^*(\mathbf{y})$ maximizes $U_{i,j}(\mathbf{z}) = V_{i,j} + \tau(b_{i,j} - A_i)$, over all alternatives $j \in \mathcal{R}_i$.

Proof of (i) Assume, by contradiction, that $j_i^*(\mathbf{z})$ is LP dominated by alternatives j and j' , i.e., $b_{i,j} < b_{i,j_i^*(\mathbf{z})} < b_{i,j'}$ and $y_{i,j} < y_{i,j_i^*(\mathbf{z})} < y_{i,j'}$, and

$$\frac{b_{i,j'} - b_{i,j_i^*(\mathbf{z})}}{w_{i,j'} - w_{i,j_i^*(\mathbf{z})}} \geq \frac{b_{i,j_i^*(\mathbf{z})} - b_{i,j}}{w_{i,j_i^*(\mathbf{z})} - w_{i,j}}.$$

From equation (8), $w_{i,j} = V_{i,j_i^*(0)} - V_{i,j}$ and thus the previous condition can be written as

$$\frac{b_{i,j'} - b_{i,j_i^*(\mathbf{z})}}{V_{i,j_i^*(\mathbf{z})} - V_{i,j'}} \geq \frac{b_{i,j_i^*(\mathbf{z})} - b_{i,j}}{V_{i,j} - V_{i,j_i^*(\mathbf{z})}}.$$

Multiplying by $\tau > 0$ on both sides and adding and subtracting A_i yields

$$\tau \frac{b_{i,j'} - b_{i,j_i^*(\mathbf{z})} - A + A}{V_{i,j_i^*(\mathbf{z})} - V_{i,j'}} \geq \tau \frac{b_{i,j_i^*(\mathbf{z})} - b_{i,j} - A + A}{V_{i,j} - V_{i,j_i^*(\mathbf{z})}}.$$

Rearranging the terms and using equation (21) yields

$$\frac{z_{i,j'} - z_{i,j_i^*(\mathbf{z})}}{V_{i,j_i^*(\mathbf{z})} - V_{i,j'}} \geq \frac{z_{i,j_i^*(\mathbf{z})} - z_{i,j}}{V_{i,j} - V_{i,j_i^*(\mathbf{z})}}.$$

Finally, using equation (1), we get

$$\frac{U_{i,j'}(\mathbf{z}) - U_{i,j_i^*(\mathbf{z})}(\mathbf{z})}{V_{i,j_i^*(\mathbf{z})} - V_{i,j'}} \geq \frac{U_{i,j_i^*(\mathbf{z})}(\mathbf{z}) - U_{i,j}(\mathbf{z})}{V_{i,j} - V_{i,j_i^*(\mathbf{z})}}.$$

Let $\alpha = V_{i,j} - V_{i,j_i^*(\mathbf{z})}$ and $\alpha' = V_{i,j_i^*(\mathbf{z})} - V_{i,j'}$. From $y_{i,j} < y_{i,j_i^*(\mathbf{z})} < y_{i,j'}$, it follows that $V_{i,j} > V_{i,j_i^*(\mathbf{z})} > V_{i,j'}$ and thus $\alpha, \alpha' > 0$. Then, we get

$$\frac{U_{i,j'}(\mathbf{z}) - U_{i,j_i^*(\mathbf{z})}(\mathbf{z})}{\alpha'} \geq \frac{U_{i,j_i^*(\mathbf{z})}(\mathbf{z}) - U_{i,j}(\mathbf{z})}{\alpha}.$$

By simple arithmetic calculation, one can see that this is equivalent to

$$\frac{\alpha U_{i,j'}(\mathbf{z}) + \alpha' U_{i,j}(\mathbf{z})}{\alpha + \alpha'} \geq U_{i,j_i^*(\mathbf{z})}(\mathbf{z}). \quad (\text{A6})$$

Equation (A6) means that the utility of $j_i^*(\mathbf{z})$ is less than or equal to the weighted average of the utility of j and j' . Two cases could then hold:

- Either j or j' is preferred to $j_i^*(\mathbf{z})$, i.e., $U_{i,j}(\mathbf{z}) > U_{i,j_i^*(\mathbf{z})}(\mathbf{z})$ or $U_{i,j'}(\mathbf{z}) > U_{i,j_i^*(\mathbf{z})}(\mathbf{z})$. This would mean that $j_i^*(\mathbf{z})$ does not maximize utility and would contradict equation (4).
- The three alternatives are equivalent, i.e., $U_{i,j}(\mathbf{z}) = U_{i,j_i^*(\mathbf{z})}(\mathbf{z}) = U_{i,j'}(\mathbf{z})$. This would also contradict equation (4) because $b_{i,j'} > b_{i,j_i^*(\mathbf{z})}$, by assumption.

Therefore, $j_i^*(\mathbf{z})$ is not a LP-dominated alternative. Clearly, $j_i^*(\mathbf{z})$ is not dominated either and thus $j_i^*(\mathbf{z}) \in \mathcal{R}_i$.

Proof of (ii) The proof of (ii) requires the following lemmas.

Lemma A *If the alternatives in \mathcal{R}_i are ordered according to equation (11), i.e., they are ordered by increasing weight, then $\tilde{\epsilon}_{i,1} > \tilde{\epsilon}_{i,2} > \dots > \tilde{\epsilon}_{i,r_i}$.*

To prove this lemma, show by contradiction that if $\tilde{\epsilon}_{i,j} \leq \tilde{\epsilon}_{i,j+1}$, then j would be LP-dominated by $j-1$ and $j+1$.

Lemma B *If $j \in \mathcal{R}_i$ is such that $\tilde{\epsilon}_{i,j} \geq 1/\tau$, then $U_{i,j-1}(\mathbf{z}) \leq U_{i,j}(\mathbf{z})$, where $j-1$ denotes the alternative which comes just before j in the ordered set \mathcal{R}_i .*

To prove this lemma, note that, using equations (8), (12) and (13), the inequality $\tilde{\epsilon}_{i,j} \geq 1/\tau$ can be written as

$$\frac{b_{i,j} - b_{i,j-1}}{V_{i,j-1} - V_{i,j}} \geq 1/\tau.$$

Multiplying by $\tau(V_{i,j} - V_{i,j-1}) > 0$, subtracting $\tau \cdot A_i$ from both sides and rearranging the terms, we get $V_{i,j-1} - \tau(A - b_{i,j-1}) \leq V_{i,j} - \tau(A - b_{i,j})$. Using equations (21) and (1) yields $U_{i,j-1}(\mathbf{z}) \leq U_{i,j}(\mathbf{z})$.

Lemma C *If $j \in \mathcal{R}_i$ is such that $\tilde{\epsilon}_{i,j} < 1/\tau$, then $U_{i,j-1}(\mathbf{z}) > U_{i,j}(\mathbf{z})$, where $j-1$ denotes the alternative which comes just before j in the ordered set \mathcal{R}_i .*

This lemma can be proved with the same reasoning as Lemma B.

Let $j \in \mathcal{R}_i$ be such that

$$\tilde{\epsilon}_{i,j} \geq 1/\tau > \tilde{\epsilon}_{i,j+1}. \quad (\text{A7})$$

Then, Lemmas A and B imply that the alternatives in the set $\{j' \in \mathcal{R}_i : j' \leq j\}$ are ordered by non-decreasing utility. Similarly, Lemmas A and C imply that the alternatives in the set $\{j' \in \mathcal{R}_i : j' \geq j\}$ are ordered by decreasing utility. Hence, alternative j , defined by equation (A7), is the alternative which maximizes the utility $U_{i,j}(\mathbf{z})$, over all alternatives in \mathcal{R}_i .

Observe that, by construction of Algorithm 1, the alternative $j_i^*(\mathbf{y})$ satisfies equation (A7). Hence, alternative $j_i^*(\mathbf{y})$ maximizes the utility $U_{i,j}(\mathbf{z})$, over all alternatives $j \in \mathcal{R}_i$. As we have shown that $j_i^*(\mathbf{z}) \in \mathcal{R}_i$, it must be that $j_i^*(\mathbf{y}) = j_i^*(\mathbf{z})$. \square

Proof of Proposition 6.1. Let $\xi = \epsilon_{i,j_i^*(0)} - \epsilon_{i,j}$. Then, ξ is the difference of two i.i.d. Gumbel's distributed random variables with scale parameter μ , and thus it is a logistic-distributed random variable with scale parameter μ (Nadarajah and Kotz 2005). Its probability density function is $f(x) = \frac{e^{x/\mu}}{\mu(e^{x/\mu} + 1)^2}$ and its cumulative distribution function is $F(x) = \frac{e^{x/\mu}}{e^{x/\mu} + 1}$. For any $z \in \mathbb{R}$, we have

$$\begin{aligned} \mathbb{E}(\xi | \xi > z) &= \frac{\int_z^\infty x f(x) dx}{1 - F(z)} \\ &= \frac{1}{1 - e^{z/\mu} / (1 + e^{z/\mu})} \int_z^\infty \frac{x e^{x/\mu}}{\mu(e^{x/\mu} + 1)^2} dx \\ &= (1 + e^{z/\mu}) \int_z^\infty \frac{x e^{x/\mu}}{\mu(e^{x/\mu} + 1)^2} dx. \end{aligned}$$

Using

$$\frac{\partial}{\partial x} \left(-\frac{x}{e^{x/\mu} + 1} \right) = \frac{x e^{x/\mu}}{\mu(e^{x/\mu} + 1)^2} - \frac{1}{e^{x/\mu} + 1},$$

and

$$\frac{\partial}{\partial x} \mu \ln(1 + e^{-x/\mu}) = \frac{-e^{-x/\mu}}{1 + e^{-x/\mu}} = \frac{-1}{e^{x/\mu} + 1},$$

we get

$$\begin{aligned} \int_z^\infty \frac{x e^{x/\mu}}{\mu(e^{x/\mu} + 1)^2} dx &= \frac{z}{e^{z/\mu} + 1} - \int_z^\infty \frac{-1}{e^{x/\mu} + 1} dx. \\ &= \frac{z}{e^{z/\mu} + 1} + \mu \ln(1 + e^{-z/\mu}). \end{aligned}$$

Finally, we obtain the required result

$$\begin{aligned} y_{i,j} &= \hat{y}_{i,j} + \mathbb{E}(\xi | \xi > -\hat{y}_{i,j}) \\ &= \hat{y}_{i,j} - \hat{y}_{i,j} + \mu(1 + e^{-\hat{y}_{i,j}/\mu}) \ln(1 + e^{\hat{y}_{i,j}/\mu}) \\ &= \mu \frac{1 + e^{\hat{y}_{i,j}/\mu}}{e^{\hat{y}_{i,j}/\mu}} \ln(1 + e^{\hat{y}_{i,j}/\mu}). \end{aligned}$$

□

Appendix B. Comparison with Proportional-Incentive Policy and Tripod

A proportional-incentive policy \mathbf{z} is a proportional tax-subsidy policy, as in (21) where only subsidies and not taxes are distributed, i.e.

$$A_i \leq b_{i,j}, \quad \forall i \in \mathcal{I}, j \in \mathcal{N}_i. \quad (\text{B1})$$

An example of proportional-incentive policy is Tripod (Araldo et al. 2019).

In this section we show that proportional-incentive policies are inefficient incentive policies, in the sense that, to achieve a certain social welfare level, they spend more incentives than needed. We call ‘inefficiency gap’ this additional incentive spent and we compute a lower bound for it in the following proposition.

Proposition B.1. *Consider a proportional-incentive policy \mathbf{z} as before. There always exists a personalized-incentive policy \mathbf{y} that is able to achieve at least the same social welfare and provides the following savings in the amount of incentive spent:*

$$Y(\mathbf{z}) - Y(\mathbf{y}) \geq \frac{1}{\tau} \cdot \sum_{i \in \mathcal{I}} (V_{i,j_i^*}(\mathbf{z}) - V_{i,j_i^*}(0)) \cdot \Delta e_{i,j_i^*}(\mathbf{z}) \equiv L(\mathbf{z})$$

where $\Delta e_{i,j} \equiv e_{i,j} - 1/\tau$ is defined as efficiency loss, $\forall i \in \mathcal{I}, j \in \mathcal{N}_i$. The quantity $L(\mathbf{z})$ defined above is a lower bound for the inefficiency gap.

Note that the *efficiency loss* quantifies the fact that proportional-incentive policies are not able to exploit the inherent efficiency $e_{i,j}$ (see Definition 4.1) of the incentivized alternative. Indeed, instead of using such an alternative-dependent efficiency, they use a single value $1/\tau$.

For the next proposition, we make use of the following definition.

Definition B.2 (Maximum Step Size and Characteristic Incremental Efficiency). Let us run Algorithm 1 with a certain budget Q and record the values $Y^{[k]}$ calculated therein, as well as the incremental efficiency of the split item $\tilde{e}_{s,t}$. The maximum step size is defined as $\gamma_Q \equiv \max_{k=1,2,\dots}(Y^{[k]} - Y^{[k-1]})$. The characteristic incremental efficiency of budget Q is defined as $\tilde{e}_Q \equiv \tilde{e}_{s,t}$.

We now compute a lower bound on the suboptimality gap of proportional-incentive policies.

Proposition B.3. *Let us consider a proportional-incentive policy \mathbf{z} , achieving a social welfare $B(\mathbf{z})$ and spending an incentive amount $Y(\mathbf{z})$. Let us denote with $B^*(Y(\mathbf{z}))$ the maximum social welfare achievable by an optimal personalized-incentive policy with that incentive amount. The following lower bound holds on the suboptimality gap:*

$$B^*(Y(\mathbf{z})) - B(\mathbf{z}) \geq \max \{0, \tilde{e}_{Y(\mathbf{z})} \cdot (L(\mathbf{z}) - 2\gamma_{Y(\mathbf{z})})\}$$

where $\gamma_{Y(\mathbf{z})}$ and $\tilde{e}_{Y(\mathbf{z})}$ are the maximum step size and the characteristic incremental efficiency, as defined in Definition B.2.

We now draw an interesting parallel with Tripod, a proportional incentive policy described in Araldo et al. (2019). In Tripod, social welfare is represented, in particular, by energy reduction. While our formulation is general and can encompass any type of social welfare (provided that the assumptions of Section 3 are valid) in our case study (Section 7) we consider CO₂ reduction. In both our case study and Tripod, individuals are travellers and alternatives are modal choices. An important aspect of Tripod is that it is dynamic, i.e., time is slotted and, in each time slot, the incentives for the individuals happening to depart in that time slot are calculated. With the Tripod policy, that we denote \mathbf{z}^{Tr} , the incentives proposed to individuals are proportional to the gain in the social indicator with respect to the default alternative, i.e.

$$z_{i,j}^{\text{Tr}} = (b_{i,j} - b_{i,j_i^*}(0))/\text{TEE}_t,$$

where the constant TEE_t is called Token Energy Efficiency and is fixed by the regulator in every time slot t . In Tripod, the incentives are distributed to individuals under a First-Come First-Served discipline, until a certain budget is depleted. Therefore, out of the entire population \mathcal{I}_t of individuals departing at time slot t only a subset $\mathcal{I}_t^{\text{Tr}}$ actually receive an incentive. In Araldo et al. (2019), the value of TEE_t is fixed empirically, with a grid search, trying several values of TEE_t in simulation and choosing the one with maximum social welfare. The calculation is based on a Model-Predictive Control (MPC) setting, where at every time slot the value of TEE_t is calculated taking into account not only the current time slot, but also a prediction of the system state (congestion, individual arrival) in the subsequent time-slots, which are called *optimization horizon*. Note that the MPC setting of Tripod allows to take into account the impact of the incentive policies on congestion, which we instead neglect, based on Assumptions 3.2 and 3.3. Therefore, for adopting our policy to a real scenario, care should be taken in checking that such Assumptions are reasonable, as we do in our case study.

Within our framework, Tripod can be defined as a proportional-incentive policy, with $\tau = 1/\text{TEE}_t$ and $A_i = b_{i,j_i^*}(0)$, applied to population $\mathcal{I}_t^{\text{Tr}}$.

As a consequence of Proposition B.1, Tripod is an inefficient incentive policy, under the assumptions of Section 3.1. In Corollary B.4, we lower-bound the additional

incentive spent with respect to the theoretical best incentive policy.

Corollary B.4. *For any Tripod incentive policy \mathbf{z}^{Tr} , there always exists a personalized-incentive policy \mathbf{z} that is able to achieve at least the same social welfare while spending less incentives. The saving in the incentives is:*

$$Y(\mathbf{z}^{Tr}) - Y(\mathbf{z}) \geq \sum_t TEE_t \cdot \sum_{i \in \mathcal{I}_t^{Tr}} (V_{i,j_i^*}(\mathbf{z}^{Tr}) - V_{i,j_i^*}(0)) \cdot \Delta e_{i,j_i^*}(\mathbf{z}^{Tr}),$$

where $\Delta e_{i,j} \equiv e_{i,j} - TEE_t$ is defined as efficiency loss, $\forall i \in \mathcal{I}_t^{Tr}, j \in \mathcal{N}_i$ and \mathcal{I}_t^{Tr} is the set of individuals getting incentives in Tripod in time-slot t .

Corollary B.4 shows that Tripod is far from minimizing the incentives needed to obtain a certain social welfare, while the policy issued by Algorithm 1 is generally close to using minimal incentives. This is confirmed by our numerical results in Section 7.6. An interpretation of the inefficiency suffered by Tripod follows.

Remark 1. Tripod uses a single value $TEE_t = 1/\tau$ to compute incentives for all alternatives j of all users $i \in \mathcal{I}_t$ and gets $1/TEE_t$ additional units of social welfare per additional unit of incentive spent. The only incentivized alternatives are the ones for which $e_{i,j} \geq TEE_t$ (otherwise the proposed incentive would not be accepted by the individual). In other words, Tripod gets always an efficiency (unit of social welfare improvement over unit of incentive spend) that is lower than the intrinsic efficiency of the incentivized alternatives. By contrast, our personalized policy always entirely exploits the intrinsic efficiency of the incentivized alternatives.

Since Tripod is a proportional-incentive policy, the same lower bound on the sub-optimality gap as in Proposition B.3 also holds, which we do not write for the sake of space.

Appendix C. Census Data Description

We now describe the census data we use.⁵ They are published by INSEE and concern the period from 2015 to 2019. The data contain observations for 7 861 201 households, representing 21 810 707 individuals (about a third of national population). Only one individual is surveyed in each household, which means that, for example, the main mode of transportation is only observed for one individual in the household. Hence, in each household, we consider only the surveyed individual.

We restrict our sample to workers living and working in the Rhône department, with a valid mode of transportation (i.e., unemployed and individuals working from home are excluded). We remove some outliers, i.e., individuals travelling more than 90 minutes, which were about 2000. The final dataset contains 221 571 individuals. The total number of alternatives is 1 092 748.

Note that census data do not represent an exhaustive sample of the population. Therefore, some categories of individuals might be over- or under-represented. To correct for such imbalances, INSEE computes a weight for each surveyed person. To compute the statistics below and to perform the multinomial regression, we use these weights.

⁵<https://www.insee.fr/fr/statistiques/4507890>

Home and Work Location. The home and work location of the individuals is reported at the city-level, except for Lyon where it is reported at the district-level. There are 275 unique home locations (an average of 812 individuals living at each location).

Mode of Transportation. The main mode of transportation used for commuting is in one of the following five categories (with mode shares in parenthesis): car (60.7%), public transit (25.1%), walking (8.8%), cycling (3.9%) and motorcycle (1.5%).

Socio-Demographic Variables. The data contain socio-demographic variables which are used to estimate a Multinomial Logit model for mode choice. The numeric variables are the age of the individual (rounded to the nearest five-year age group, mean is 38.5 years old) and the number of cars owned per employed individuals in the household (mean is 0.84). The categorical variables are the gender of the individual (50.7% of individuals are men) and the occupation of the individual (using INSEE nomenclature; the most frequent category is “employee” with 24.9% of individuals).

Appendix D. Computation of Travel Times

For any individual, the origin point of her trips is set to the town hall of the city where she lives and the destination point is set to the town hall of the city where she works (for district-level home and workplace, the town hall of the district is used). The coordinates of the town halls are retrieved from OpenStreetMap.

Travel time of each mode is set to the travel time of the fastest path, within that mode, which connects the two locations, computed from the open-source routing engine GraphHopper. The road network for pedestrians, bicycles, motorcycles and cars is retrieved from OpenStreetMap data. It is assumed that there is no congestion. For public transit trips, the fastest path is computed using public transit timetables, retrieved from open-data GTFS files. The departure time is assumed to be at 8 a.m. on a weekday.

Note that, for some individuals, no path can be found to travel by public transit from their origin to their destination (16 161 individuals, representing 10.87% of total sample weight). For these individuals, we exclude public transit from their choice set.

Some individuals are living and working in the same city (61 497 individuals, representing 26.61% of total sample weight). For these individuals, travel times are computed by supposing that trip distance is equal to the radius of the city (assuming cities are circular) and that speed is equal to the average speed of intercity trips.

Figure D1 shows the distribution of travel times in the population, for each mode of transportation. Except for public transit trips, most trips last less than 30 minutes.

Appendix E. Simulating Utilities

In the Multinomial Logit model, the utility of individual i with mode of transportation j is $V_{i,j} = \hat{V}_{i,j} + \epsilon_{i,j}$, where $\hat{V}_{i,j}$ is the deterministic part of the utility, which depends on the individual- and alternative-specific exogenous variables, and $\epsilon_{i,j}$ is a random variable with standard Gumbel distribution.

The deterministic part $\hat{V}_{i,j}$ are computed from the estimates of the Multinomial




Figure D1. Distribution of travel times in the population (before the policy).

Logit model. From the data, we know the alternative $j_i^*(0)$ chosen by any individual i so we must have

$$V_{i,j_i^*(0)} > V_{i,j}, \quad \forall j \neq j_i^*(0). \quad (\text{E1})$$

To simulate draws of standard Gumbel variables conditional on equation (E1), we use the rejection sampling method, i.e., we draw values from the standard Gumbel distribution until the constraint of equation (E1) is satisfied.