Online Companion to "Collaborative Vehicle-to-grid Operations in Frequency Regulation Markets"

Appendix A: The Persistency Model as a Generalized Choice Model

We provide a summary of the marginal moment model (MMM) (Natarajan et al. 2009), which we shall utilize in our modeling. We consider a generic zero-one linear program in the form of:

$$\max_{\boldsymbol{x}\in\boldsymbol{\Lambda}} \sum_{a\in A} v_a x_a, \quad \text{where } \boldsymbol{\Lambda} \equiv \left\{ \boldsymbol{x} \in \{0,1\}^{|A|} : \sum_{a\in A} b_{na} x_a = f_n, \quad \text{for } n \in N \right\}. \tag{23}$$

We assume that the continuous relaxation of (23) is tight, such that the optimal objective value does not change if $\boldsymbol{\Lambda}$ is replaced with its convex hull, $conv(\boldsymbol{\Lambda})$. Examples that satisfy this assumption include network-based problems (e.g., shortest path, longest path on acylic networks, assignment problems) with totally unimodular constraint matrices. We are interested in the behavior of the optimal solution $\boldsymbol{x}$ given random values of the coefficients $\boldsymbol{v}$. That is, given a random variable $\tilde{\boldsymbol{v}}$ of which $\boldsymbol{v}$ is one realization, we are interested in the corresponding optimal solutions as a random variable $\tilde{\boldsymbol{x}}$. The metric of particular interest, given that $\tilde{\boldsymbol{x}} \in \{0,1\}^{|A|}$, is the probability that each component will equal one in an optimal solution, i.e., $P(\tilde{x}_a = 1)$ for each $a \in A$.

As exact analysis has proved difficult, the literature has focused on finding approximations and bounds that can be computed efficiently, often through convex optimization formulations. For the MMM in Natarajan et al. (2009), only the first two marginal moments of $\tilde{\boldsymbol{v}}$, but not the precise joint distribution, are known. That is, it is given that $\tilde{v}_a \in \mathbb{R}$, $E[\tilde{v}_a] = \mu_a$ and $E[\tilde{v}_a^2] = \sigma_a^2$. Natarajan et al. (2009) prove a convex optimization formulation that yields a tight upper bound on the expected value of (23) that is valid under all distributions with these given moments, and exactly attained by an *extremal* distribution. Furthermore, the optimal solution to the optimization problem yield the probabilities that the decision variables are equal to one at optimality ($P(\tilde{x}_a = 1)$) under the extremal distribution, i.e., the persistency values. As Lemma 2 in Natarajan et al. (2009) suggests:

THEOREM 1. *A tight upper bound on the expected value of* (23) *under the MMM can be computed by solving the following concave maximization problem:*

$$\max_{\mathbf{y}\in conv(\boldsymbol{\Lambda})} \sum_{a\in A} \left(\mu_a y_a + \sigma_a \sqrt{y_a(1-y_a)} \right). \tag{24}$$

Furthermore, the optimal value of y_a gives the associated persistency value $P(\tilde{x}_a = 1)$ for each $a \in A$.

Appendix B: Proofs and Background Results

Proof of Proposition 1

The Lagrangian dual of (5) is given by

$$\min_{\boldsymbol{\rho}\in\mathbb{R}^{|N|}} \max_{\mathbf{y}\in[0,1]^{|A|}} \sum_{a\in A}\left[\left(\mu_a - \sum_{n\in N} b_{na}\rho_n\right) y_a + \sigma_a\sqrt{y_a(1-y_a)}\right] + \sum_{t=1}^{T}\left[\left(\mu_t + s_t - \sum_{n\in N} b_{nt}\rho_n\right) y_t + \sigma_t\sqrt{y_t(1-y_t)}\right] + \sum_{n\in N} f_n\rho_n.$$

The optimal solution to the inner maximization problem is given by:

$$y_a^* = \frac{1}{2}\left(1 + \frac{\lambda_a}{\sqrt{\lambda_a^2+\sigma_a^2}}\right), \quad \text{for } a\in A, \tag{25}$$

$$y_t^* = \frac{1}{2}\left(1 + \frac{\lambda_t}{\sqrt{\lambda_t^2+\sigma_t^2}}\right), \quad \text{for } t=1,\ldots,T, \tag{26}$$

where λ_a and λ_t are defined according to (10) and (11). Note that $y_a^*\in[0,1]$ for all $a\in A$. Then, the dual can be rewritten as (27):

$$\Gamma(\mathbf{s}) = \min_{\boldsymbol{\rho}\in\mathbb{R}^{|N|}} \frac{1}{2}\sum_{a\in A}\left(\lambda_a + \sqrt{\lambda_a^2+\sigma_a^2}\right) + \frac{1}{2}\sum_{t=1}^{T}\left(\lambda_t + \sqrt{\lambda_t^2+\sigma_t^2}\right) + \sum_{n\in N} f_n\rho_n, \tag{27}$$

Strong duality holds due to concavity of (5) and the existence of a strictly feasible solution.

Finally, constraint (9) imposes the strong duality condition, i.e., primal (maximization) objective is greater than or equal to dual (minimization) objective. □

Definition of Deviation Measures and Proof of Proposition 2

Definition 1. For a zero-mean random variable $\tilde{z}$, the forward and backward deviations are defined as follows, respectively:

$$\sigma^f(\tilde{z}) = \sup_{\theta>0}\{\frac{1}{\theta}\sqrt{2\ln(E[\exp(\theta\tilde{z})])}\}$$
$$\sigma^b(\tilde{z}) = \sup_{\theta>0}\{\frac{1}{\theta}\sqrt{2\ln(E[\exp(-\theta\tilde{z})])}\}$$

The deviation measures are significant due to some of their properties, (e.g., Chen et al. 2007, Proposition 2):

Proposition 4. *Let $\sigma, \sigma^f, \sigma^b$ be the standard, forward and backward deviations, respectively, for a zero-mean random variable $\tilde{z}$. Then:*

1. *$\sigma^f \geq \sigma$ and $\sigma^b \geq \sigma$.*
2. *If $\tilde{z}$ follows a Normal distribution, $\sigma^f = \sigma^b = \sigma$.*
3. *For any $\theta\geq 0$, $P(\tilde{z}\leq\theta\sigma^f)\leq\exp(-\theta^2/2)$ and $P(\tilde{z}\geq-\theta\sigma^b)\leq\exp(-\theta^2/2)$.*

4. *For any* $\theta \geq 0$, $\ln E[\exp(\theta\tilde{z})] \leq \theta^2(\sigma^f)^2/2$ *and* $\ln E[\exp(-\theta\tilde{z})] \leq \theta^2(\sigma^b)^2/2$.

Making use of Proposition 4, it is possible to bound the CVaR function as follows:

PROPOSITION 5. *Consider a random vector* $\tilde{\mathbf{z}} \in \mathbb{R}^I$ *whose components are mutually independent, and have zero means and finite forward and backward deviations. Then,*

$$\phi_{1-\eta}\left(\alpha_0 + \sum_{i=1}^{I} \alpha_i \tilde{z}_i\right) \leq \alpha_0 + \sqrt{-2\ln\eta}\sqrt{\sum_{i=1}^{I} u_i^2} \tag{28}$$

$$\text{where } u_i = \max\{\sigma_i^f \alpha_i, -\sigma_i^b \alpha_i\}$$

$$\phi_{1-\eta}\left(\alpha_0 + \sum_{i=1}^{I} \alpha_i \tilde{z}_i\right) \leq \alpha_0 + \frac{1-\eta}{\eta}\sqrt{-2\ln(1-\eta)}\sqrt{\sum_{i=1}^{I} v_i^2} \tag{29}$$

$$\text{where } v_i = \max\{-\sigma_i^f \alpha_i, \sigma_i^b \alpha_i\}.$$

Proof of Proposition 2

We apply Proposition 5 to bound (12). Recall that $\tilde{K}_t$ follows a binomial distribution. Invoking the Normal approximation of the binomial and Proposition 4.2, the forward and backward deviations of $\tilde{K}_t$ are both equal to its standard deviation, i.e., $\sqrt{My_t(1-y_t)}$. Then, (28) becomes:

$$\bar{r}_t C_t - \psi M y_t + \sqrt{-2\ln\eta}\sqrt{\hat{u}_t^2 + u_t^2} \leq 0$$
$$\text{where } \hat{u}_t = \psi\sqrt{My_t(1-y_t)}, u_t = \sigma_t^f C_t,$$

which yields (13). Similarly, (29) becomes:

$$\bar{r}_t C_t - \psi M y_t + \frac{1-\eta}{\eta}\sqrt{-2\ln(1-\eta)}\sqrt{\hat{v}_t^2 + v_t^2} \leq 0$$
$$\text{where } \hat{v}_t = \psi\sqrt{My_t(1-y_t)}, v_t = \sigma_t^b C_t,$$

which yields (14).

Proof of Proposition 3

Constraints (17) and (18) jointly enforce (9). Furthermore, since (9) enforces the strong duality condition that holds if and only if the associated variables (including $\mathbf{y}$) constitute an optimal primal-dual pair, constraint (17) must hold at strict equality. This implies that (13) and (14) are equivalent to (15) and (16).

Appendix C: Linearization of Bilinear Terms

To linearize the bilinear terms $s_t y_t$, we provide a disjunctive formulation as follows. First, suppose for each period t, the rebate s_t is chosen out of a set of given candidate rebate values $\{\hat{s}_{th} : h \in H\}$. Let $z_{th} = 1$ if rebate level $\hat{s}_{th}$ is selected, and 0 if not. This can be imposed by the following constraints:

$$\begin{aligned} s_t &= \sum_{h \in H} \hat{s}_{th} z_{th} \\ \sum_{h \in H} z_{th} &= 1; \quad z_{th} \in \{0,1\} \text{ for } h \in H. \end{aligned}$$

The last set of constraints require that one and only one of z_{th} can be non-zero for each t. They can be rewritten as special ordered set (SOS) constraints to take advantage of more efficient branch-and-bound algorithms under commercial solver packages such as Gurobi and CPLEX.

Then, we define the new variable u_t to replace the bilinear term $s_t y_t$. This relationship can be imposed by a set of disjunctive constraints. In particular, define variables u_{th} and y_{th} indexed by $h \in H$. Then, the following constraints ensure that $u_t = s_t y_t$:

$$\begin{aligned} y_{th} &\le z_{th}, \quad \text{for } h \in H \\ u_{th} &\le z_{th} \hat{s}_{th}, \quad \text{for } h \in H \\ u_{th} &= y_{th} \hat{s}_{th}, \quad \text{for } h \in H \\ \sum_{h \in H} y_{th} &= y_t \\ \sum_{h \in H} u_{th} &= u_t. \end{aligned}$$

Appendix D: Complete Model Formulation and Parameter Values

The full MISOCP formulation is given below.

$$\max \ \sum_{t=1}^{T} (\bar{p}_t C_t - u_t M)$$

$$\text{s.t.} \ \sum_{a\in A} \left(\mu_a y_a + \sigma_a \sqrt{y_a(1-y_a)}\right) + \sum_{t=1}^{T} (\mu_t y_t + u_t + \sigma_t w_t) \geq$$

$$\frac{1}{2}\sum_{a\in A} \left(\lambda_a + \sqrt{\lambda_a^2 + \sigma_a^2}\right) + \frac{1}{2}\sum_{t=1}^{T} \left(\lambda_t + \sqrt{\lambda_t^2 + \sigma_t^2}\right) + \sum_{n\in N} f_n \rho_n,$$

$$\sum_{a\in A} b_{na} y_a + \sum_{t=1}^{T} b_{nt} y_t = f_n, \quad \text{for } n \in N,$$

$$\lambda_a = \mu_a - \sum_{n\in N} b_{na}\rho_n, \quad \text{for } a \in A,$$

$$\lambda_t = \mu_t + s_t - \sum_{n\in N} b_{nt}\rho_n, \quad \text{for } t = 1, \cdots, T,$$

$$\bar{r}_t C_t - \psi M y_t + \sqrt{-2\ln\eta}\sqrt{\psi^2 M w_t^2 + (\sigma_t^f C_t)^2} \leq 0, \quad \text{for } t = 1, \cdots, T,$$

$$w_t^2 + y_t^2 \leq y_t, \quad \text{for } t = 1, \cdots, T,$$

$$s_t = \sum_{h\in H} \hat{s}_{th} z_{th}, \quad \text{for } t = 1, \cdots, T,$$

$$\sum_{h\in H} z_{th} = 1, \quad \text{for } t = 1, \ldots, T, \text{ and } h \in H,$$

$$y_{th} \leq z_{th}, \quad \text{for } t = 1, \ldots, T, \text{ and } h \in H,$$

$$u_{th} \leq z_{th}\hat{s}_{th}, \quad \text{for } t = 1, \ldots, T, \text{ and } h \in H,$$

$$u_{th} = y_{th}\hat{s}_{th}, \quad \text{for } t = 1, \ldots, T, \text{ and } h \in H,$$

$$\sum_{h\in H} y_{th} = y_t, \quad \text{for } t = 1, \ldots, T,$$

$$\sum_{h\in H} u_{th} = u_t, \quad \text{for } t = 1, \ldots, T,$$

$$z_{th} \in \{0,1\}, \quad \text{for } t = 1, \ldots, T, \text{ and } h \in H, \quad \text{[indicator for selecting rebate level } h \text{ in period } t\text{]}$$

$$y_{th}, u_{th} \geq 0, \quad \text{for } t = 1, \ldots, T, \text{ and } h \in H, \quad \text{[auxiliary variables for linearization, see App. C]}$$

$$s_t, u_t \geq 0, \quad \text{for } t = 1, \ldots, T, \quad \text{[auxiliary variables for linearization, see App. C]}$$

$$C_t \geq 0, \quad \text{for } t = 1, \ldots, T, \quad \text{[capacity bid in period } t\text{]}$$

$$0 \leq y_a \leq 1, \quad \text{for } a \in A, \quad 0 \leq y_t \leq 1, \quad \text{for } 1 \leq t \leq T. \quad \text{[persistency values]}$$

Finally, the parameter values for the numerical experiments are summarized below.

Table 8 Parameter Values Used in the Computational Experiments

Parameter	Value	Rationale
M	100 to 400	The fleet size is assumed to be 200 in the base case. This is varied in Section 4.3.1.
$\bar{p}_t$	[12.00, 11.78, 9.25, 9.49, 11.62, 23.49, 32.58, 31.67, 20.72, 20.10, 30.00, 24.76, 27.97, 28.55, 31.91, 43.97, 51.28, 72.72, 65.74, 30.78, 29.33, 29.88, 26.13, 28.99]	The average (over May 2018) hourly market clearing prices are obtained from the PJM frequency regulation market (PJM 2021a). These are varied by a scaling factor in the sensitivity analysis in Section 4.3.2.
$\bar{r}_t$	[0.642, 0.623, 0.600, 0.613, 0.620, 0.619, 0.621, 0.600, 0.625, 0.638, 0.645, 0.623, 0.643, 0.638, 0.629, 0.608, 0.608, 0.530, 0.546, 0.535, 0.533, 0.552, 0.644, 0.614]	The hourly value of $\tilde{r}_t$, the minimum power required to fulfill the minimum mileage requirement for a 1 MW bid, is obtained based on the PJM AGC signals for each day in May 2018 PJM (2021c). $\bar{r}_t$ gives the hourly average over the month.
σ_t^f	[0.046, 0.052, 0.061, 0.051, 0.050, 0.035, 0.072, 0.058, 0.063, 0.032, 0.041, 0.043, 0.048, 0.034, 0.063, 0.043, 0.055, 0.099, 0.086, 0.082, 0.064, 0.079, 0.046, 0.056]	The forward deviations of $\tilde{r}_t$, calibrated following Chen et al. (2007).
ψ	7.2 kW	The charger's power rating is assumed to be the standard 7.2 kW output for most Level-2 chargers. This is varied in Section 4.3.3.
η	0.05	Chance constraint for fulfilling the mileage requirement guarantees no more than 5% violation probability.
$\boldsymbol{\mu}$	Values omitted for brevity	The average utility values for activity and idle arcs are calibrated with the inverse optimization procedure discussed in Section 3.4
$\boldsymbol{\sigma}$	Values omitted for brevity	The standard deviations of utility values for activity and idle arcs are calibrated with the inverse optimization procedure discussed in Section 3.4

Appendix E: Computational Enhancements

E.1: Valid Inequalities

In Section 3, we have reformulated the platform's pricing problem as a MISOCP (see Appendix D), which can be readily solved by commercial solvers such as Gurobi and CPLEX. To further enhance computational performance, we derive a class of valid inequalities based on the structure of the scheduling network to strengthen the formulation. First, we show that the optimal persistency values ($\mathbf{y}$) exhibits monotonicity with respect to the rebates ($\mathbf{s}$) .

LEMMA 1. *The optimal flow y_t in the persistency model is non-decreasing in $s_{t'}$ for $t' \neq t \in \{1, \ldots, T\}$.*

Lemma 1 suggests that increasing rebates in any period has a complementary effect of incentivizing drivers to plug in *other* periods. This result arises from the structure of the scheduling network (Figure 3), where the charging arcs are aligned as complements (Granot and Veinott 1985) of one another. Intuitively, this holds because idling the EV in some period t' (to take advantage of increased rebates) precludes any activity that covers the interval $[t, t']$ (or $[t', t]$), and thus

increases the probability of the EV idling in period t. This result shall be useful in developing valid inequalities to tighten the formulation of the consolidated problem.

Lemma 1 helps tighten the formulation through improved bounds on the $\mathbf{y}$ variables. Denote by $\underline{y}_t$ the optimal flow for charging arc $t \in \{1,\dots,T\}$ in the pricing optimization model with extra constraints $s_{t'} = 0$, for $t' \in \{1,\dots,T\}\backslash\{t\}$. Similarly, denote by $\bar{y}_t$ the optimal flow for charging arc $t \in \{1,\dots,T\}$ in the pricing optimization model with the extra constraints $s_{t'} = \bar{s}$, for $t' \in \{1,\dots,T\}\backslash\{t\}$. Then, the following holds:

PROPOSITION 6. *The optimal flow for arc $t \in \{1,\dots,T\}$ satisfies $\underline{y}_t \le y_t \le \bar{y}_t$.*

Proposition 6 provides valid bounds for the persistency values. Recall that the terms $s_t y_t$ are linearized using a mixed integer programming formulation (see Appendix C) by discretizing the range of the rebate variables s_t. During the branch-and-bound process, tightening the range of y_t will help cut off fractional solutions and reduce the search space. We will discuss numerical experiments in Appendix E that validates the computational improvements.

Proof of Lemma 1

Interpreting the persistency values $\mathbf{y}$ as flows in the scheduling network (Figure 3), the persistency model (5) is in the form of a network flow problem with a concave objective in the maximization sense (i.e., convex cost by reversing the sign).

The time network is biconnected (Granot and Veinott 1985) because every pair of distinct arcs is contained in a simple cycle. Moreover, in this network, any two charging arcs $t \neq t' \in \{1,\dots,T\}$ are complements, since every simple cycle containing both t and t' would have both oriented in the same direction (see Granot and Veinott 1985, for the definition of complements and substitutes). Further, $-\Gamma(\mathbf{y},\mathbf{s})$ is convex and continuous for any $\mathbf{s}$, and is submodular (referred to as "subadditive" in Granot and Veinott (1985)) for $t \in \{1,\dots,T\}$.

These conditions imply that the optimal flow y_t is non-decreasing in $s_{t'}$ for any $t' \neq t \in \{1,\dots,T\}$ (by Granot and Veinott 1985, Theorem 10). □

Proof of Proposition 6

Consider $t \in \{1,\dots,T\}$, the optimal rebates $s_{t'}$ for $t' \neq t \in \{1,\dots,T\}$ satisfy $0 \le s_{t'} \le \bar{s}$. Therefore, by Lemma 1, the optimal flow y_t for the persistency subproblem satisfies $\underline{y}_t \le y_t \le \bar{y}_t$. □

E.2: Computational Performance

In this subsection, we illustrate the computational performance of our model and the effectiveness of the valid inequalities derived in Appendix E.1. Specifically, we generate a set of instances with randomly perturbed driver utility profiles, and report the solution times and MISOCP optimality gaps (if not solved to optimality within 30 minutes), with and without adding the valid inequalities (from Proposition 6 in Appendix E.1).

To generate different instances, we first apply agglomerative hierarchical clustering on the travel patterns of residents. Specifically, respondents in the travel survey sample can be clustered into two distinct groups by their travel patterns, as shown in Figure 8.

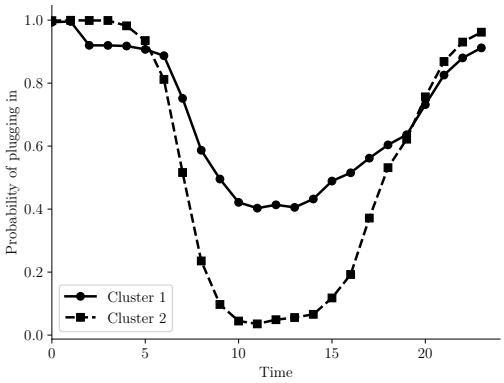

Figure 8 Travel Pattern Clusters in CHTS Sample

Cluster 1 (plotted with a solid line), consisting of about 60% of the sample, spends more time out of home compared to those in Cluster 2 (dashed line). These individuals also tend to depart from home earlier in the morning and return slightly earlier in the afternoon. This clustering is very intuitive: Cluster 1 represents individuals who commute to work following the regular 9-5 work schedule; Cluster 2 consists of those who either have more flexible commuting schedules/modes or do not commute to work.

We generate a set of test problem instances by varying the composition of the two clusters in the sample. For each composition (row in Table 9), we generate 50 instances and report the average computational performance. We generate a sample of daily schedules for 5,000 simulated drivers for each instance, and assign each driver to Clusters 1 or 2 with probabilities consistent with the setting; Then, for each hour, we simulate the driver's status following the conditional probabilities under each cluster (Figure 8). After collecting the simulated sample of drivers for each instance, we re-estimate the utility parameters $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ following the inverse optimization procedure in Section 3.4.

The experiments are conducted by using the Gurobi 9.1.0 solver with Python 3.6.8 on a 64-bit Linux workstation with an Intel Xeon CPU (20 cores) at 2.6GHz and 256 GB RAM. When solving

each instance, we set a limit of 1,800s on the solution time. The valid inequalities (Proposition 6 in Appendix E.1) are added to the formulation with a callback function in Gurobi, i.e., are added whenever they are violated in an interim solution during branch and bound. Since not all instances are solved within the time limit, we also report the percentage of instances successfully solved to optimality in each case. We calculate the average computing time only for the solved instances and record the remaining optimality gaps for the unsolved instances.

Table 9 Computational Performance

Composition		without valid inequalities			with valid inequalities		
Cluster 1	Cluster 2	Time(s)	Gap	% Solved	Time(s)	Gap	% Solved
0.4	0.6	174.02	-	100%	42.46	-	100%
0.45	0.55	202.68	-	100%	58.31	-	100%
0.5	0.5	298.45	2.38%	94%	54.78	-	100%
0.55	0.45	490.88	2.23%	86%	102.80	-	100%
0.6	0.4	676.23	2.41%	62%	171.98	2.13%	98%
0.65	0.35	758.29	2.46%	40%	330.58	2.20%	86%
0.7	0.3	994.54	2.60%	18%	473.09	2.29%	64%

The results are summarized in Table 9. Across the settings tested, we find that adding the valid inequalities significantly improve both the percentage of instances solved optimally and the solution times for solved instances (more than halved). This suggests that the valid inequalities are highly effective in enhancing computational performance. Each instance involves solving an MISOCP with 2209 continuous variables and 384 binary variables, and 672 quadratic constraints and 24 special ordered set (SOS) constraints. In the vast majority of cases, the problem can be solved optimally within 10 minutes.

Interestingly, the computational difficulty depends on the representations of Clusters 1 and 2 in the sample of drivers. Recall that drivers in Cluster 2 already stay at home with a large probability without rebates. Therefore, when Cluster 2 takes a larger proportion, the platform will offer less (often close to zero) rebates, and the corresponding optimization problem becomes easier to solve. Yet, for all instances, the tightened formulation is computationally tractable.

Appendix F: Maximum Likelihood Estimation of Utility Parameters

In this appendix, we illustrate that our inverse optimization formulation (20) can also be used as a heuristic solution to a maximum likelihood estimation (MLE) approach to estimate the utility parameters $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$.

Mishra et al. (2014) propose a procedure to estimate the parameters of the MMM model in a discrete choice context using MLE. While our problem is more complex (due to the underlying network structure of the choices), we may adapt their procedure as discussed below. Similar to the case of binary classification considered in Mishra et al. (2014), we can define the likelihood function for the discrete choice model as follows, given sample persistency values $\hat{\mathbf{y}}$:

$$\prod_{a\in\mathcal{A}} \Pr\{\text{Choosing arc } a|\boldsymbol{\mu},\boldsymbol{\sigma}\}^{\hat{y}_a} \prod_{t=1}^{T} \Pr\{\text{Choosing arc } t|\boldsymbol{\mu},\boldsymbol{\sigma}\}^{\hat{y}_t}.$$

Note that the probability of drivers choosing a specific arc has been calculated by (25) and (26), then we can formulate the maximum log-likelihood problem as follows:

$$\begin{aligned}
\max_{\boldsymbol{\lambda},\boldsymbol{\rho},\boldsymbol{\mu},\boldsymbol{\sigma}} \quad & \sum_{a\in A} \hat{y}_a \ln\left(\frac{1}{2}+\frac{\lambda_a}{2\sqrt{\lambda_a^2+\sigma_a^2}}\right) + \sum_{t=1}^{T} \hat{y}_t \ln\left(\frac{1}{2}+\frac{\lambda_t}{2\sqrt{\lambda_t^2+\sigma_t^2}}\right) \\
\text{s.t.} \quad & \lambda_a = \mu_a - \sum_{n\in N} b_{na}\rho_n, \quad \forall a\in A, \\
& \lambda_t = \mu_t + s_t - \sum_{n\in N} b_{nt}\rho_n, \quad \forall t=1,\dots,T, \\
& \frac{1}{2}\sum_{a\in A} b_{na}\left(1+\frac{\lambda_a}{\sqrt{\lambda_a^2+\sigma_a^2}}\right) + \frac{1}{2}\sum_{t=1}^{T} b_{nt}\left(1+\frac{\lambda_t}{\sqrt{\lambda_t^2+\sigma_t^2}}\right) = f_n, \quad \forall n\in N.
\end{aligned}$$

This MLE problem is nonconvex. To solve this MLE problem, we may convexify the formulation with the variable substitution $z_a = \left(\frac{1}{2}+\frac{\lambda_a}{2\sqrt{\lambda_a^2+\sigma_a^2}}\right)$. This allows the problem to be decomposed into two steps. First, we may solve for $\mathbf{z}$:

$$\begin{aligned}
\max_{\mathbf{z}} \quad & \sum_{a\in A} \hat{y}_a \ln z_a + \sum_{t=1}^{T} \hat{y}_t \ln z_t, \\
\text{s.t.} \quad & z_a \in (0,1] \text{ for } a\in A, \quad z_t \in (0,1] \text{ for } t=1,\dots,T, \\
& \mathbf{z} \in conv(\boldsymbol{\Lambda}).
\end{aligned}$$

Then, we solve for the remaining variables $(\boldsymbol{\lambda},\boldsymbol{\mu},\boldsymbol{\sigma},\boldsymbol{\rho})$ that yields the resulting persistency values z_a. The latter step results in a subproblem that is identical to the inverse optimization problem as in Proposition 3. If the inverse optimization problem is solved with zero objective (duality gap), the resulting solution exactly solves the MLE problem. Otherwise, if there is a small but non-zero duality gap, the solution may serve as a heuristic MLE estimate (since the original MLE problem is nonconvex).

Therefore, a (heuristic) MLE estimation is equivalent to our original inverse optimization approach; but with the observed persistency values given by the MLE estimates z_a, instead of sample average $\hat{y}_a$ values in the inverse optimization formulation.

To assess the performances, we estimate the utility parameters $(\boldsymbol{\mu}, \boldsymbol{\sigma})$ under our original inverse optimization approach and the heuristic MLE approach, and re-solve the persistency model (24) to produce predictions of the persistency values accordingly as a result of the estimated utility parameters $(\boldsymbol{\mu}, \boldsymbol{\sigma})$. Let the two sets of predicted persistency values be denoted $\boldsymbol{y}_{inv}$ and $\boldsymbol{y}_{MLE}$, respectively. We then compare these against the observed persistency values in the CHTS data, both in sample $(\boldsymbol{y}_{test})$ and out of sample $(\boldsymbol{y}_{test})$ using 10-fold cross validation.

We calculate the average cosine similarity and the root mean squared error (RMSE) to test the in-sample and out-of-sample performances of both methods, as reported below.

Table 10 Prediction accuracy of persistency values using inverse optimization and MLE estimates of the utility parameters

	Cosine similarity $\frac{\boldsymbol{y}_{true}\boldsymbol{y}_{est}}{\|\boldsymbol{y}_{true}\|_2\|\boldsymbol{y}_{est}\|_2}$		RMSE $\sqrt{\frac{\|\boldsymbol{y}_{true}-\boldsymbol{y}_{est}\|_2^2}{N}}$	
	in-sample	out-of-sample	in-sample	out-of-sample
Inverse Optimization	99.97%	97.77%	6.6e-3	0.04
MLE	96.74%	94.91%	0.04	0.06

Table 10 shows that both methods demonstrate strong out-of-sample performance, with the inverse optimization method slightly outperforming the MLE method. Consequently, we choose to retain the direct inverse optimization method in our main text. In our estimation, the MLE approach does not lead to better predictions or richer insights unfortunately. This is due to the fact that the CHTS dataset only gives a sample of observed driving scheduling decisions, whereas no treatment (i.e., rebates in our model) or exogenous shocks are observed. In the future, for example, if the V2G platform performs experiments to test drivers' responses to different rebate schedules as a treatment, employing the MLE approach will be fruitful over the sample average approach.

Appendix G: Glossary of Application-Specific Terminology

Term	Explanation
PJM grid	PJM is a regional transmission organization (RTO) in the United States. It is part of the Eastern Interconnection grid operating an electric transmission system serving all or parts of Delaware, Illinois, Indiana, Kentucky, Maryland, Michigan, New Jersey, North Carolina, Ohio, Pennsylvania, Tennessee, Virginia, West Virginia, and the District of Columbia.
Resource	In power grids, "resources" refer to energy generation and storage facilities such as power plants, solar or wind farms, grid battery storage, V2G systems, etc.
Frequency regulation	Frequency regulation is a type of ancilliary service that the grid (e.g., PJM) procures from energy resources such as power plants to help stabilize the grid's frequency (e.g., at 50 Hz).
Dispatch	Dispatch refers to the amount of energy (in MWh) that an energy resource supplies to the grid.
Power (Capacity)	Power refers the rate at which energy is transferred (between an energy resource and the grid), in MW. The capacity of a resource is the maximum power that it is capable of dispatching.
Regulation up/down	When demand for energy exceeds (resp., drops below) generation temporarily in the grid, the frequency drops (resp. rises). To maintain constant frequency, the system operator (e.g., PJM) requests regulation resources to dispatch (withdraw) power to the grid to counteract the change in frequency. This is known as regulation up (resp., down).
Automatic Generation Control (AGC)	In real time, the system operator communicates with different power generators and energy resources to control the rate at which they exchange (dispatch or withdraw) energy with the grid. This is done electronically through an AGC signal.
Mileage	A measure of accuracy to which a resource (e.g., storage battery) participating in frequency regulation is able to track the AGC signal by dispatching power in response to it.
Mileage-based performance	The PJM grid requires energy resources participating in frequency regulation to be able to exchange power with the grid by tracking the AGC signal accurately, i.e., the mileage must exceed a minimum standard.

Table 11 Glossary of application-specific terms