# Algorithmic Precision and Human Decision: A Study of Interactive Optimization for School Schedules

Arthur Delarue, Zhen Lian, Sébastien Martin *

**Abstract**

In collaboration with the San Francisco Unified School District (SFUSD), this paper introduces an interactive optimization framework to tackle complex school scheduling challenges. The choice of school start and end times is an optimization challenge, as schedules influence the district's transportation system, and limiting the associated costs is a computationally difficult combinatorial problem. However, it is also a policy challenge, as transportation costs are far from the only consequence of school schedule changes. Policymakers need time and knowledge to balance these considerations and reach a consensus carefully; past implementations have failed because of policy issues despite state-of-the-art optimization approaches. We first motivate our approach with a micro-foundation model of the interplay between policymakers and researchers, arguing that limiting their dependency is key. Building on these insights, we propose a framework that includes (1) a fast algorithm capable of solving the school schedule problem that compares favorably to the literature and (2) an interactive optimization approach that leverages this speed to allow policymakers to explore a variety of solutions in a transparent and efficient way, facilitating the policy decisionmaking process. The framework led to the first optimization-driven school start time changes in the US, updating the schedule of all 133 schools in SFUSD in 2021 with annual transportation savings exceeding $5 million. A comprehensive survey of approximately 27,000 parents and staff in 2022 provides evidence of the approach's effectiveness.

**Key words:** Optimization, public policy, interactive, decision support, impact, education

*Arthur Delarue: H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, USA; arthur.delarue@isye.gatech.edu. Zhen Lian: Yale School of Management, New Haven, CT 06525; zhen.lian@yale.edu. Sébastien Martin: Kellogg School of Management, Northwestern University, Evanston, IL, USA; sebastien.martin@kellogg.northwestern.edu.

# 1 Introduction

Teenagers in the United States start school too early, according to the American Academy of Pediatrics (Owens et al. 2014) and the American Medical Association 2016. Shifts in the body's internal rhythm during puberty make it harder to fall asleep early without decreasing the total amount of time teens must sleep (Crowley et al. 2018). As a result, high school bells ringing before 8:30 can harm students' mental and physical health. Many school districts have sought to address these concerns in recent years by moving high school start times later, and studies have shown later start times lead to more sleep and better health (Meltzer et al. 2021). Yet most U.S. teenagers still begin school too early.

The status quo of too-early start times persists primarily because school districts face considerable operational obstacles in aligning school start times with medical guidelines. For budget reasons, school districts spread the start and end times of different schools, such that a school bus dropping off teenagers at 7:30 can go on to pick up younger students starting later in the morning instead of waiting at the depot until the afternoon. Thus, school start and end times significantly impact transportation spending. In addition to these operational implications, start and end times have a material influence on family schedules, as parents must plan around school pickups and dropoffs. The schedules also impact the students' extracurriculars, school staffing, and even public transportation. Therefore, no option will satisfy everyone, and the choice of start times is a policy challenge.

This interplay between a thorny policy problem and an intractable operational problem makes changing school start times particularly challenging. Despite a state-of-the-art optimization approach and a unanimous school committee vote, a proposal by Bertsimas, Delarue, Eger, et al. 2020 in Boston was not implemented after staunch opposition from some parent groups. In contrast, our work uses interactive optimization to give policymakers more control over the outcome and overcome the policy hurdles. We designed this approach in collaboration with the San Francisco Unified School District (SFUSD), the seventh-largest district in California, with a total enrollment of over 50,000 across more than 120 schools (SFUSD 2021). Our work led to successfully implementing a new school schedule in San Francisco in the 2021-2022 school year, which is to our knowledge the first optimization-driven school start time change in the United States. This change saves five million dollars annually to this day in transportation costs and was well received by parents and staff.

Our work leverages the example of SFUSD to provide a principled approach to the school start time scheduling problem. To provide micro-foundations for our interactive optimization approach, we first introduce a multi-period model of policymakers' and researchers' collaboration progress. Inspired by our experience with SFUSD, this model highlights the inter-dependency between research and policy, where each party's progress depends on the other. Researchers can solve challenging optimization problems, but policymakers' expertise is required to align their models with the complex reality. However, policymakers cannot reach a consensus without knowing what is feasible and thus rely on the researchers' work to explore the solution space. Asymptotic analysis of this setting reveals a strong bottleneck effect: in the best case, progress is driven by the slowest party. Our results are also prescriptive: reducing the dependency between research and policy is beneficial, even if it lowers the research speed. We interpret this result as a motivation for interactive optimization: researchers can design interactive tools that allow policymakers to explore the solution space freely, and this is worth it even if building the tools is time-consuming.

A main contribution of this work is the introduction of such interactive tools. A typical approach to solving multi-objective problems such as school start time choice is to generate Pareto curves to help policymakers control the optimization process, as in Bertsimas, Delarue, Eger, et al. 2020. For example, SFUSD wanted to minimize transportation costs and changes to the old schedule, and it is hard to balance these objectives ex-ante without observing solutions. This *objective prioritization* approach was only the first step in our work: in high-dimensional problems such as start time choice, the near-optimal solution space can still be vast and varied even for a specific objective prioritization. These degrees of freedom are typically resolved almost arbitrarily by optimization solvers but can instead be leveraged by policymakers to create a more inclusive process and overcome policy challenges. To enable this more transparent approach, we generate thousands of near-optimal solutions and design interactive tools that policymakers can use to explore this space efficiently. In San Francisco, these tools led to an unanticipated phenomenon of *objective discovery*: policymakers realized that much more could be achieved than their initial objectives. So, they quickly iterated in just a few weeks and chose a schedule that alleviated many staff and parents' concerns. To confirm this, we ran a large-scale survey involving 24,000 families and over 3,500 staff members. The survey shows overall satisfaction with the new start times and some evidence that the interactive optimization process added value to the stakeholders compared to a default implementation.

However, this interactive approach is computationally intensive, requiring thousands of high-quality solutions in short timelines. At the same time, the school start time optimization problem is notoriously hard due to the difficulty of optimizing transportation costs. We introduce a tractable yet accurate integer programming approach that approximates transportation costs and is orders of magnitude faster than state-of-the-art alternatives. Large-scale routing simulations show that this approximation performs well compared to various baselines. This formulation is paired with a sampling technique to generate solutions for interactive optimization, balancing between the sub-optimality of the solutions and their diversity, measured by an entropy metric. This approach extends the interactive optimization literature to the case where some objectives are unknown *ex-ante*.

### 1.1 Related Literature

**School start times and transportation.** The negative impact of too-early school starts on teenagers' health has been widely documented, leading to worse academic performance (Carrell, Maghakian, and West 2011), car accidents (Danner and Phillips 2008), and other health problems (Crowley et al. 2018). The American Academy of Pediatrics and the American Medical Association both recommend later start times for high schools. In school districts that have successfully overhauled start and end times, Goldin et al. (2020), Widome et al. (2020) and Meltzer et al. (2021) empirically show that moving high school start times later does lead to more sleep for teenage students. Yet changing start times remains a thorny policy question because of its complex influence on school operations, particularly transportation.

Roughly half of U.S. public school students ride a yellow school bus each day. Park and B. I. Kim 2010 and Ellegood et al. 2020 survey school transportation from an operations research perspective. The school bus routing problem is typically decomposed into three subproblems: (i) stop assignment, or clustering students into combined pickup locations; (ii) bus routing, or combining stops associated with the same school into "runs"; (iii) bus scheduling, or taking advantage of staggered school start and end times to build bus itineraries that serve several runs for different schools in succession. This last subproblem has received the most attention, starting with a classical discrete-time network flow formulation from Swersey and Ballard 1984, where start and end times are fixed. Desrosiers et al. 1986 jointly optimized bus and school schedules using an alternating minimization heuristic, while Zeng, Chopra, and Smilowitz 2022 incorporated school start times as a decision variable in a joint formulation of bus and school scheduling. Both works simplify the Swersey and Ballard 1984 formulation by ignoring transitions between bus routes, and only focusing on minimizing the maximum number of active runs at any given time.

Start and end times are usually seen as just another lever to reduce transportation cost; additional objectives are not considered, with two exceptions. Banerjee and Smilowitz 2019 consider the formulation from Zeng, Chopra, and Smilowitz 2022, and add a second objective which minimizes the maximum change from the status quo. Bertsimas, Delarue, and Martin 2019 propose a multi-objective optimization formulation for start time selection, which relies on a different approximation of transportation costs. Their approach requires solving quadratic assignment problems, providing great modeling flexibility at the expense of tractability. We seek to combine the best of both works: we develop a highly tractable optimization formulation, a simplification of Banerjee and Smilowitz 2019, while expanding Bertsimas, Delarue, and Martin 2019's focus on multi-objective exploration capabilities.

**Humans, algorithms, and interactive decision support.** The political complexity of start time change requires policymakers to directly examine many options and gather feedback from stakeholders (Owens et al. 2014; Bertsimas, Delarue, Eger, et al. 2020). The interplay between experts and models in decision-making is an area of recent interest in the operations management literature (Davis et al. 2022). While data-driven models can sift through many more possibilities than humans, the biases caused by modeling simplification can lead to human mistrust (i.e., algorithm aversion, see Dietvorst, Simmons, and Massey 2015). An effective approach to restore trust is to grant human experts more control over the model and transparency over its workings (Dietvorst, Simmons, and Massey 2018; Bolton and Katok 2018; Buell, T. Kim, and Tsay 2017).

Increasing modeling control by the end user is a key objective in *interactive optimization*. The typical multi-objective optimization approach of generating entire Pareto frontiers is not practical when there are more than two or three objectives. Interactive optimization approaches present a Pareto-optimal solution to decision-makers, then move to a preferred Pareto-optimal solution based on the feedback received. For continuous, convex problems, the solution can be updated using a variety of optimization techniques, including the Frank-Wolfe method (Geoffrion, Dyer, and Feinberg 1972), the simplex method (Zionts and Wallenius 1983), or projected gradients (Yang 1999). The literature is much sparser when the decision problem is combinatorial, with a focus on metaheuristics (Teghem, Tuyttens, and Ulungu 2000; Barbati, Corrente, and Greco 2024). Recent reviews by Miettinen, Ruiz, and Wierzbicki 2008 and Bandaru, Ng, and Deb 2017 identify interactive discrete optimization as an opportunity area. All these methods also make the crucial assumption that objective functions are known explicitly *ex ante* — an assumption that does not hold in the complex setting of school start times.

Indeed, Zechman, Giacomoni, and Shafiee 2013 define "wicked" problems as problems that defy conventional optimization modeling due to complex human decision factors. One approach is "modeling to generate alternatives" (Brill, Chang, and Hopkins 1982), which provides policymakers with multiple different "good-enough" solutions. This literature focuses on three key questions: (i) how to generate diverse high-quality solutions, (ii) how to modulate the size and diversity of the solution set, and (iii) how to design tools to enable its exploration. The first question has received the most attention, with approaches including fractional programming (Trapp and Konrad 2015), specialized cuts (Voll et al. 2015), branch-and-bound (Danna et al. 2007), and sequential optimization (Brill, Chang, and Hopkins 1982; Greistorfer et al. 2008). Our approach is most similar to Chang, Brill Jr, and Hopkins 1982, where random linear objectives are sampled while deviations from the original objectives are constrained. The second question has received less attention; many of the methods above present decision-makers with just a handful of solutions, which can be validated one by one. Considering more solutions requires new algorithms (Xin et al. 2022) and data structures (Serra and Hooker 2020). Almost all existing methods measure solution diversity using notions of pairwise distance, which lack interpretability. We generate thousands of near-optimal solutions and measure diversity using a simple entropy metric. We thus refine the analysis of Chang, Brill Jr, and Hopkins 1982 by quantifying the tradeoff between suboptimality and diversity.

Finally, the development of interactive tools to explore near-optimal solutions is an open area of research. The closest approach is by Xin et al. 2022, who non-interactively report summary statistics on a "Rashomon set" of near-optimal decision trees. Of course, decision support tools in general have a long history of success in the operations research literature. Caro et al. 2004 present a GIS-based tool which allows policy makers to tweak school attendance zones, while Chu, Keskinocak, and Villarreal 2020 propose an Excel-based tool which allows transportation planners to make small changes to school bus routes. Ahuja, Jha, and Liu 2007 describe an interactive tool to allow railroad network designers to improve the solutions produced by a large-scale optimization model. Gershenfeld 2015 develops a ticket pricing model and describes obtaining executive buy-in via a series of interactive demonstrations. In a more general context, Olavson and Fry 2008 discuss the challenges of designing useful spreadsheet-based decision tools for real applications. In a way, this paper bridges the gap between decision support tools and modeling to generate alternatives.

Section 2 uses a mathematical model based on our interactions with SFUSD to illustrate the interactions between policymakers and researchers and explain why interactive optimization can be useful. Section 3 then presents our school start time optimization formulation, the backbone of our interactive optimization approach. We also use large-scale simulations in Boston and SFUSD to argue that our tractability gains have a limited impact on the quality of the solutions. Section 4 presents our interactive optimization framework and the associated user interfaces, and Section 5 details its implementation in SFUSD. In particular, we present the results of a post-implementation survey of almost 28,000 SFUSD families and staff to provide evidence of the benefits of our approach.

## 2 A model for research-policy collaboration

Choosing school schedules represents both a complex combinatorial optimization challenge due to its vehicle routing implications and a significant policy issue, as it involves finding objectives that a broad range of stakeholders can agree on. Researchers and policymakers also depend on each other, and we illustrate this interplay with a model based on our work with SFUSD that motivates our interactive optimization approach.

### 2.1 Motivation: research and policy errors

**Optimization error.** Given a choice of policy $\boldsymbol{\lambda} \in \Lambda$, researchers must solve:

$$\begin{aligned} \text{minimize} \quad & c(\mathbf{x}, \boldsymbol{\lambda}), \\ \text{s.t.} \quad & \mathbf{x} \in \mathcal{X}, \end{aligned}$$

where $\mathbf{x} \in \mathcal{X}$ represents a feasible solution (e.g., a start and end time for each school), and the policy $\boldsymbol{\lambda}$ parametrizes the cost function: $c(\mathbf{x}, \boldsymbol{\lambda})$ is the cost of the solution $\mathbf{x}$ given the policy $\boldsymbol{\lambda}$. For example, if the school district's policy $\boldsymbol{\lambda}$ is to minimize the cost of the school bus system, then $c(\mathbf{x}, \boldsymbol{\lambda})$ could represent the expected transportation costs given a schedule $\mathbf{x}$. We measure the quality of the researchers' progress with the *optimization error* of their solution $\mathbf{x}$:

$$\varepsilon_{\text{opt}}(\mathbf{x}, \boldsymbol{\lambda}) = c(\mathbf{x}, \boldsymbol{\lambda}) - c(\mathbf{x}^*(\boldsymbol{\lambda}), \boldsymbol{\lambda}), \tag{1}$$

where $\mathbf{x}^*(\boldsymbol{\lambda}) = \arg\min_{\mathbf{x}\in\mathcal{X}} c(\mathbf{x}, \boldsymbol{\lambda})$ is the optimal solution (assumed unique for clarity of exposition) for a particular choice of policy $\boldsymbol{\lambda}$. The optimization error can arise from modeling simplifications or computational limitations. For example, when minimizing transportation costs, researchers may approximate the routing problem and obtain suboptimal solutions.

**Policy error.** Each policy $\boldsymbol{\lambda} \in \Lambda$ represents a choice of utility loss function $c(\cdot, \boldsymbol{\lambda})$ for the policymakers, capturing their objectives (e.g., minimizing transportation costs, satisfying parent/staff preferences, improving student health...) and their relative importance. Some policies could involve constraints (e.g., *all schools must start after 8 am*), which can be formulated as infinite-cost objectives. Reaching a consensus on $\boldsymbol{\lambda}$ often requires lengthy discussions and community consultations. Because time is often limited and consensus hard to achieve, there is no guarantee that the research-policy collaboration will lead to perfect policy decisions. To represent this issue, we assume that there is a "best" choice of policy $\boldsymbol{\lambda}^{\text{true}} \in \Lambda$, typically hard to find, such that $c(\cdot, \boldsymbol{\lambda}^{\text{true}})$ best represents the stakeholders' utility loss and $\mathbf{x}^*(\boldsymbol{\lambda}^{\text{true}})$ the corresponding best solution. The gap between a policy $\boldsymbol{\lambda}$ chosen by policymakers and the optimal one can be characterized by the *policy error*:

$$\varepsilon_{\text{pol}}(\boldsymbol{\lambda}) = c(\mathbf{x}^*(\boldsymbol{\lambda}), \boldsymbol{\lambda}^{\text{true}}) - c(\mathbf{x}^*(\boldsymbol{\lambda}^{\text{true}}), \boldsymbol{\lambda}^{\text{true}}). \tag{2}$$

The first term is the cost incurred if researchers perfectly optimize an imperfect policy. That is, given a policy $\boldsymbol{\lambda}$, we evaluate the cost of its optimal solution $\mathbf{x}^*(\boldsymbol{\lambda})$ with respect to the best policy $\boldsymbol{\lambda}^{\text{true}}$. The second term is simply the baseline, i.e., the lowest cost that can be achieved given $\boldsymbol{\lambda}^{\text{true}}$. Therefore, the policy error measures the true utility loss induced by choosing an imperfect policy if there is no optimization error.

**Policy-research collaboration and interactive optimization.** The policy-research collaboration aims to find a solution that minimizes the true stakeholder utility loss. Formally, the utility loss of a solution $\mathbf{x}$ is represented by the *total error*,

$$\varepsilon_{\text{tot}}(\mathbf{x}) = c(\mathbf{x}, \boldsymbol{\lambda}^{\text{true}}) - c(\mathbf{x}^*(\boldsymbol{\lambda}^{\text{true}}), \boldsymbol{\lambda}^{\text{true}}), \tag{3}$$

which represents the utility loss of $\mathbf{x}$ with respect to the true objective $\boldsymbol{\lambda}^{\text{true}}$. The total error is deeply related to the policy and optimization errors.

**Proposition 1.** *For any policy $\boldsymbol{\lambda}$, the total error can be decomposed as*

$$\varepsilon_{\text{tot}}(\mathbf{x}) = \varepsilon_{\text{pol}}(\boldsymbol{\lambda}) + \varepsilon_{\text{opt}}(\mathbf{x}, \boldsymbol{\lambda}) + \delta(\mathbf{x}, \boldsymbol{\lambda}), \tag{4}$$

*where $\delta(\mathbf{x}, \boldsymbol{\lambda}) = c(\mathbf{x}, \boldsymbol{\lambda}^{true}) - c(\mathbf{x}, \boldsymbol{\lambda})$. Furthermore,*

- *If* $\varepsilon_{\text{pol}}(\boldsymbol{\lambda}) = 0$ *(no policy error),* $\varepsilon_{\text{tot}}(\mathbf{x}) = \varepsilon_{\text{opt}}(\mathbf{x}, \boldsymbol{\lambda}^{\text{true}})$ *(optimization error remains).*
- *If* $\varepsilon_{\text{opt}}(\mathbf{x}, \boldsymbol{\lambda}) = 0$ *(no optimization error),* $\varepsilon_{\text{tot}}(\mathbf{x}) = \varepsilon_{\text{pol}}(\boldsymbol{\lambda})$ *(policy error remains).*
- *If* $\varepsilon_{\text{pol}}(\boldsymbol{\lambda}) = \varepsilon_{\text{opt}}(\mathbf{x}, \boldsymbol{\lambda}) = 0$*, then* $\varepsilon_{\text{tot}}(\mathbf{x}) = 0$.

Proposition 1 shows that the total error is not exactly the sum of policy and optimization errors, as captured by $\delta(\mathbf{x}, \boldsymbol{\lambda})$. This is because it's possible to have a wrong solution (high $\varepsilon_{\text{opt}}$) to a wrong policy (high $\varepsilon_{\text{pol}}$) that is accidentally good for the true policy (low $\varepsilon_{\text{tot}}$). This is unlikely to happen in practice. And if it does, it's also unlikely to be identified as a good solution, as researchers are sensitive to optimization error and policymakers to policy error. The proposition shows that a good way to minimize the total error is for both parties to focus on eliminating the errors they control.

Proposition 1 illustrates a common way to organize a policy-research collaboration and minimize the total error: first, policymakers consult with stakeholders to fully specify their true goals (minimizing policy error); second, researchers model the corresponding optimization problem and design an algorithm to minimize the optimization error relative to the true policy. Sometimes, however, these processes cannot be conducted independently because policy progress depends on research work. With SFUSD, we witnessed two sources of policy dependence on research: *objective prioritization* and *objective discovery*.

Objective prioritization is typical in multi-objective optimization: districts must balance important objectives like costs and preferences. Prioritizing these objectives ex-ante is challenging for policymakers, who need the help of researchers to model the objectives and showcase Pareto frontiers and specific solutions that illustrate the tradeoffs between objectives. Objective discovery was particularly important in our experience with SFUSD. Our original task was to implement later high school starts and reduce transportation costs. As we started to produce results, the analysis we shared and the results we provided helped policymakers realize many other objectives were also achievable (e.g., minimizing change, helping align staff schedules across schools). Our work helped refine the policy decisions (another instance of policy progress depending on research progress); it also meant our original models were imperfect and needed revisions.

Hence, successfully minimizing the total error $\varepsilon_{\text{tot}}$ relies on a complex interplay between the policy and research work. An immediate approach, illustrated in Figure 1, is for the policymakers and the researchers to meet and share their progress regularly. Researchers present their analysis to the policymakers, which informs the policy consensus process, while the policymakers provide feedback on the researchers' models and solutions to better align them with their updated policy objectives. However, this process involves many iterations and significant communication and may fail under tight time constraints. Instead, this work uses an interactive optimization approach (also in Figure 1), where researchers build tools that policymakers can use directly. These tools aim to facilitate the process of objective discovery, reducing the dependency of the policymakers on the researchers. We

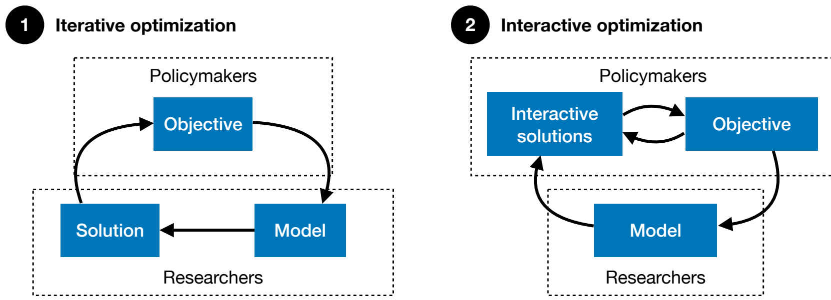

Figure 1: Two iterative collaboration methods for researchers and policymakers.

next formalize this intuition, introducing a parsimonious model that captures the essence of the dynamics we observed in our collaboration with SFUSD.

## 2.2 Collaboration model and interactive optimization

Proposition 1 illustrates the importance of jointly considering the policy error and the optimization error. We cannot model the interplay of $\varepsilon_{\text{tot}}$, $\varepsilon_{\text{opt}}$, and $\varepsilon_{\text{pol}}$ in full generality, as they are highly complex and domain-dependent. Instead, we highlight the dynamics of the research and policy collaboration in a more tractable and intuitive way.

We consider a multi-period model with $N$ periods. In each period (or iteration), researchers and policymakers make progress and reduce their respective errors. More precisely, we let $\varepsilon_{\text{pol}}^n \geq 0$ designate the policy error at iteration $n \geq 0$, while $\varepsilon_{\text{opt}}^n$ designates the optimization error. We assume that $\varepsilon_{\text{pol}}^0 = \varepsilon_{\text{opt}}^0 = 1$. The objective of the collaboration is to make the total error $\varepsilon_{\text{tot}}^n$ as close to zero as possible. In line with the discussion of Proposition 1, we assume that $\varepsilon_{\text{tot}}^n = \varepsilon_{\text{pol}}^n + \varepsilon_{\text{opt}}^n$. Importantly, policy and research are interdependent, and researchers and policymakers meet at the end of each period and communicate. The errors in period $n+1$ thus depend on both research and policy progress in period $n$. To introduce the model, we begin with two simplified cases that highlight key aspects of the dynamics.

First, consider the case where the policy error remains constant, $\varepsilon_{\text{pol}}^n = p$ for all $n$. This scenario represents a situation where policymakers make no progress on policy discussions, leaving researchers solely responsible for reducing the optimization error. In this case, we assume that the optimization error evolves as:

$$\varepsilon_{\text{opt}}^{n+1} = \varepsilon_{\text{opt}}^n (1 - \rho). \tag{5}$$

Here, the optimization error exhibits geometric decay at a rate of $1 - \rho$, and $\rho$ can be interpreted as the speed of research. The parameter $\rho$ captures diminishing marginal returns of research work: as progress accumulates, achieving further improvements becomes increasingly challenging.

Next, consider the opposite case with a constant optimization error, $\varepsilon_{\text{opt}}^n = o$ for all $n$. For example, this could model a scenario where researchers do their best to find solutions to the various policy requests, but are not able to find perfect optimization models due to tractability issues, leading to a constant optimization error $o$ regardless of the policy. In this case, there should be a limit to how much policy progress can be made (no matter how much policymakers work), given the optimization limitations and the various dependencies of policy on research we previously discussed. Hence, we assume that the policy error evolves as:

$$\varepsilon_{\text{pol}}^{n+1} = \gamma o + (\varepsilon_{\text{pol}}^n - \gamma o)(1 - \pi). \tag{6}$$

Here, $\varepsilon_{\text{pol}}^n$ converges geometrically to $\gamma o$ at a rate $1 - \pi$, where $\pi$ can be interpreted as the speed of policy progress. The parameter $\gamma$ captures the dependency of policy progress on research progress. When $\gamma = 0$, policy can progress without any research, and the policy error converges to zero at a rate of $(1 - \pi)$. Conversely, when $\gamma > 0$, policy progress depends on research and the convergence is lower-bounded by $\gamma o$ which means that policy would need improved research (lower $o$) to progress further.[1]

Having established these insights, we now introduce the full model, where both policy and research are updated at the same time:

$$\varepsilon_{\text{opt}}^{n+1} = \varepsilon_{\text{opt}}^n(1 - \rho) + (\varepsilon_{\text{pol}}^n - \varepsilon_{\text{pol}}^{n+1}), \tag{7a}$$

$$\varepsilon_{\text{pol}}^{n+1} = \gamma\varepsilon_{\text{opt}}^n + (\varepsilon_{\text{pol}}^n - \gamma\varepsilon_{\text{opt}}^n)(1 - \pi). \tag{7b}$$

We depict these dynamics in Figure 2. First, note that Eq. (7a) is equivalent to Eq. (5) if the policy error is constant, and Eq. (7b) is equivalent to Eq. (6). The only change we added to the error dynamics is the additional term $(\varepsilon_{\text{pol}}^n - \varepsilon_{\text{pol}}^{n+1})$ to the evolution of the optimization error. Indeed, when the policy is updated from $\varepsilon_{\text{pol}}^n$ to $\varepsilon_{\text{pol}}^{n+1}$, additional modeling/optimization work is required to account for the new policy goals. A simple way to illustrate this dynamic is to assume that the reduction in policy error $\varepsilon_{\text{pol}}^n - \varepsilon_{\text{pol}}^{n+1}$ is added to the optimization error. This choice allows our model to behave nicely: consider the case where the "research speed" $\rho$ is 0, that is, researchers do no work at all. Then, Eq. (7a) can be rearranged as: $\varepsilon_{\text{opt}}^{n+1} + \varepsilon_{\text{pol}}^{n+1} = \varepsilon_{\text{opt}}^n + \varepsilon_{\text{pol}}^n$, which means that $\varepsilon_{\text{tot}}^{n+1} = \varepsilon_{\text{tot}}^n$. In that case, even if policy can progress and $\varepsilon_{\text{pol}}^n$ decreases, researchers do not update their model/algorithms, and no better solution is found - the total error is constant over time. This is consistent with the definition of the total error in Eq. (3), which only depends on the current solution and should not change if researchers do not update their model/algorithms. Together, these equations capture the dynamic interplay between research and policy based

[1] The asymmetry between Eq. (5) and Eq. (6) is intentional: when policymakers rely on solutions generated by optimization methods provided by researchers (captured by $\gamma > 0$), the optimization error directly influences how good the policy can be. In contrast, a fixed policy objective, even if it is "incorrect", does not prevent researchers from optimizing it well.

on our collaboration with SFUSD. Research continually strives to align with evolving policy objectives, while policy progress is inherently reliant on the insights and outputs generated by research.

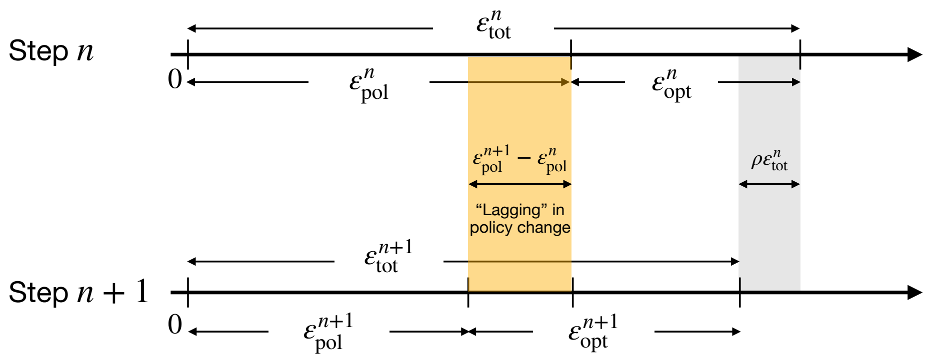


Figure 2: Illustration of model dynamics

*Note.* This figure illustrates the evolution of $\varepsilon^n_{\text{opt}}$, $\varepsilon^n_{\text{pol}}$, and $\varepsilon^n_{\text{tot}}$ over time. The orange bar represents the additional term $(\varepsilon^n_{\text{pol}} - \varepsilon^{n+1}_{\text{pol}})$ added to the optimization error at step $n+1$, caused by the optimization process catching up with the updated policy. The gray bar indicates the change in total error ($\varepsilon_{\text{tot}}$) from step $n$ to step $n+1$. When $\rho = 0$ (researchers do not make updates), the gray bar shrinks to zero, showing that the total error remains constant, despite potential progress in policy.

**Insights.** It is easy to see that the total error converges to zero when the number of iterations $N$ increases, provided $\gamma < 1$ and $\rho, \pi > 0$. The following lemma characterizes this convergence.

**Lemma 1.** *If $\gamma < 1$ and $\rho, \pi > 0$, then $\varepsilon^N_{\text{tot}} = \Theta((1-\alpha)^N)$ in the limit $N \to \infty$, with*

$$\alpha = \begin{cases} \frac{\rho+\pi(1+\gamma)}{2}\left(1 - \sqrt{1 - \frac{4\pi\rho}{(\rho+\pi(1+\gamma))^2}}\right), & \text{if } \gamma > 0, \\ \pi, & \text{if } \gamma = 0. \end{cases} \tag{8}$$

In the limit, both policy and optimization errors (and hence the total error) converge geometrically at rate $1-\alpha$ — in other words, $\alpha$ characterizes the speed of the overall collaboration progress, just as $\rho$ and $\pi$ capture the speeds of research and policy progress. A larger $\alpha$ represents a faster asymptotic convergence. The relationship between $\alpha$ and the underlying parameters $\rho$, $\pi$, and $\gamma$ is illustrated in Figure 11 of the appendix. While it seems complex, it turns out that the expression for $\alpha$ is very close to the minimum of $\rho$ and $\pi$, especially when $\gamma$ is low.

**Theorem 1** (Bottleneck effect.). *$\alpha$ increases with the policy speed $\pi$ and the research speed $\rho$, and decreases with the research dependency $\gamma$. Furthermore,*

$$\lim_{\gamma \to 0} \alpha = \min(\pi, \rho).$$

This result establishes a key role of the dependency $\gamma$. While it is immediate that $\gamma$ slows the collaboration process, it is less obvious that the convergence rate of the total error $\varepsilon^n_{\text{tot}}$ is at least as slow as the slowest speed of policy or research as long as $\gamma > 0$. Together with Lemma 1, Theorem 1 shows that there is a discontinuity: when $\gamma = 0$ (no dependency), the convergence rate for the policy error is $(1 - \pi)$. However, when $\gamma > 0$, no matter how small the dependency, progress is at least as slow as the slowest of research or policy. This bottleneck effect formalizes our intuition that problems like school start time choice are not simply optimization problems but a complex collaboration between research and policy. Designing policy-research interactions should take into account the importance of this dependency. For example, the following result shows that researchers should set up interactive tools that reduce the policymaker's dependency on them, even if it slows down the research process.

**Theorem 2.** *Let $k = \frac{\rho}{\pi}$ designate the speed ratio between research and policy. As long as the research is fast enough ($k \geq \gamma/3$), we have that*

$$-\frac{\partial \alpha}{\partial k} - \frac{\partial \alpha}{\partial \gamma} \geq 0,$$

*i.e., decreasing the research-to-policy speed ratio and the research dependency by the same infinitesimal amount improves the overall speed of convergence.*

We view this result as a formalization of interactive optimization's role in the success of the SFUSD project. Interactive optimization reduces the policy dependency on research ($\gamma$) by giving the policymakers some automated tools to help with objective discovery and prioritization without waiting for additional researcher results. Of course, researchers need additional time to set up these tools and update them when the policymakers clarify their policy goals, reducing their speed $\rho$. Theorem 2 shows that this tradeoff is beneficial, even when research is already substantially slower than policy (small $k$), and even when $\gamma$ is already small. This mechanism relies on the ability to design interactive optimization tools that can help policymakers with the challenges of objective prioritization and discovery without the need for researcher input. The next two sections describe our approach in the context of start time selection.

# 3 Optimization model

Interactive optimization should be as responsive as possible (as discussed in Section 1). Thus, our first challenge is to design a start-time optimization approach that can quickly generate solutions with limited loss of accuracy (avoiding high optimization error).

### 3.1 Optimizing school start times

We consider a set of $S$ schools indexed by $1,\dots,S$, and a set of $T$ possible start times indexed by $1,\dots,T$. For instance, SFUSD considered $T = 3$ possible start times: 7:50 AM, 8:40 AM, or 9:30 AM; we denote the $t$-th start time option by $h_t$. The binary decision variable $x_{s,t}$ takes the value 1 if school $s$ starts at time $t$, and 0 otherwise. We write the set of feasible start time solutions using a classic assignment formulation as

$$\mathcal{X} = \left\{ \mathbf{x} \in \{0,1\}^{S\times T} : \sum_{t=1}^{T} x_{s,t} = 1 \ \forall s = 1,\dots,S \right\},$$

with a constraint ensuring that each school starts at exactly one start time. A challenge in choosing the formulation in the context of interactive optimization is to have the flexibility to model a variety of objectives — indeed, as described in Section 2, policymakers may discover new objectives during the collaboration. Without loss of generality, we will focus on the following two objectives that were most important to SFUSD:

$$\begin{aligned} \text{minimize} \quad & \lambda_{\text{cost}} f_{\text{cost}}(\mathbf{x}) + \lambda_{\text{change}} f_{\text{change}}(\mathbf{x}), \\ \text{s.t.} \quad & \mathbf{x} \in \mathcal{X}, \end{aligned} \tag{9}$$

where $f_{\text{cost}}(\mathbf{x})$ represents the school bus costs of a choice of start times $\mathbf{x}$, and $f_{\text{change}}(\mathbf{x})$ is the total change in start times compared to the status quo. $\lambda_{\text{cost}}$ and $\lambda_{\text{change}}$ are their corresponding prioritization weights. These two objectives are both nonlinear, but the latter is trivially linearizable when defined as

$$f_{\text{change}}(\mathbf{x}) = \sum_{s=1}^{S} \left| \sum_{t=1}^{T} h_t x_{s,t} - o_s \right|,$$

where $o_s$ is the original start time of school $s$. Clearly, $f_{\text{change}}$ can be linearized using at most $S$ auxiliary variables. Choosing a tractable $f_{\text{cost}}$ is more challenging and will be addressed in the next part. We could also add other objectives that districts may care about to Equation (9), or additional constraints: as discussed in Section 5, SFUSD wanted all nearby elementary schools to start simultaneously. As any constraints can be modeled as an objective with a very high cost if the constraint is unmet, we will only refer to all of the district's objectives and constraints as "objectives" to simplify the exposition. Overall, the approach we describe in this paper does not depend on the choice of objectives but requires that they be added in a tractable way to the formulation (Appendix B.1 of the appendix discusses another example). The only one that was challenging for us to simplify, and which is central to most school districts is $f_{\text{cost}}$, which we discuss next.

### 3.2 A linear model of transportation cost

As described in Section 1, routing school buses optimally is difficult even when start times are fixed, and computational complexity increases significantly once start times become decision variables. We develop an approximation of transportation costs that can be linearized without too many additional decision variables.

**Active runs.** We describe our input data as follows. Each school $s$ is associated with a set of morning and afternoon *runs*. Each run corresponds to a sequence of stops visited by a single bus to transport students to and from school $s$. Morning runs (AM runs) for school $s$, indexed by $r \in \{1, \dots, R_s^{\mathrm{AM}}\}$, begin at a student's bus stop and end at school $s$, possibly picking up additional students along the way. Afternoon runs (PM runs) for school $s$, indexed by $r \in \{1, \dots, R_s^{\mathrm{PM}}\}$, begin at school $s$ and end at a student's bus stop, possibly dropping off additional students along the way. We denote by $\delta_{s,r}$ the duration of run $r$ for school $s$. If school $s$ starts at time $h_t$, morning run $r$ begins at time $h_t - \delta_{s,r}$ and ends at time $h_t$.[2] Correspondingly, afternoon run $r'$ begins at time $h_t + \ell_s$ and ends at time $h_t + \ell_s + \delta_{s,r}$, where $\ell_s$ designates the length of the instructional day at school $s$, in minutes. We say that a run is *active* at time $\theta$ if $\theta$ lies between the start time and end time of the run. In other words, AM run $r$ for school $s$ is active at time $\theta$ if $h_t - \delta_{s,r} \leq \theta \leq h_t$, and PM run $r'$ for school $s$ is active at time $\theta$ if $h_t + \ell_s \leq \theta \leq h_t + \ell_s + \delta_{s,r}$. The total number of active AM runs and PM runs at time $\theta$ for school $s$ when it starts at time $h_t$ is given by:

$$N_{s,t,\theta}^{\mathrm{AM}} = \sum_{r=1}^{R_s^{\mathrm{AM}}} \mathbb{1}\left(h_t - \delta_{s,r} \leq \theta \leq h_t\right), \tag{10a}$$

$$N_{s,t,\theta}^{\mathrm{PM}} = \sum_{r=1}^{R_s^{\mathrm{PM}}} \mathbb{1}\left(h_t + \ell_s \leq \theta \leq h_t + \ell_s + \delta_{s,r}\right), \tag{10b}$$

The tensors $\mathbf{N}^{\mathrm{AM}}$ and $\mathbf{N}^{\mathrm{PM}}$ can easily be pre-computed for a discrete set of values of $\theta$ (e.g. every 5 minutes), which we denote by $\Theta$. We can visualize the computation for one particular school in the morning in Figure 3.

**Transportation cost formulation.** In school transportation, the primary driver of cost is the size of the bus fleet, so minimizing transportation costs typically means minimizing the total number of buses. A useful approximation (and valid lower bound, see Desrosiers et al. 1986; Zeng, Chopra, and Smilowitz 2022) of the number of buses needed to serve all runs is the maximum number of active runs at any given time. This approach disregards

[2] We assume here that the length of run $\delta_{s,r}$ does not depend on the start time of the school $h_t$. This may not always be a good approximation if travel times are not stationary, but our formulation can easily be modified in that case without any loss of tractability.

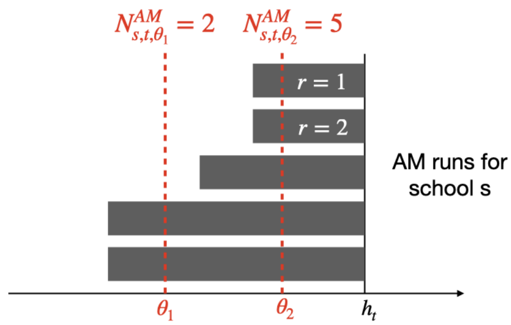


Figure 3: Illustration of the morning transportation cost for school $s$ when starting at time $h_t$

*Note.* The $x$-axis represents time, and the dark rectangles represent active runs, each with a particular length in time. The dashed red lines indicate two particular times $\theta_1$ and $\theta_2$, respectively with 2 and 5 active runs.

the transfer time of the buses between consecutive runs, and the fact that runs can be re-optimized given start time choices. However, optimal bus schedules typically have short transfers in large and dense school districts such as SFUSD, and we will see that the second stage routing decision, important to minimize costs, has a limited impact on the optimal choice of start times. This approximate transportation cost can be expressed as

$$f_{\text{cost}}(\mathbf{x}) = \max_{\theta \in \Theta} \left( \left( N^{\text{AM}}_{s,t,\theta} + N^{\text{PM}}_{s,t,\theta} \right) x_{s,t} \right), \tag{11}$$

which can be trivially linearized using just one auxiliary decision variable since it is the maximum of $|\Theta|$ linear functions in $\mathbf{x}$ (see Eq. (15) in the appendix for details). Indeed, Banerjee and Smilowitz 2019 also use the number of active runs as a proxy for transportation costs, but their method allows each run to arrive within a specific time window before the start time, introducing potential variations in the number of active runs for a given set of start times. Our method can thus be viewed as a simplification which collapses the time window to the start time. This simplification enables pre-computation of the active run tensor, which significantly improves efficiency. Furthermore, as we discuss below, the simplification leads to minimal loss of optimality.

**Extensions.** This formulation is versatile, and we improve it to capture important considerations. For instance, runs change with year-over-year enrollment shifts while start times stay fixed. Therefore, we minimize the average transportation costs under different scenarios, each corresponding to a past year. This can be done by adding another dimension to the tensor $\mathbf{N}$, denoting as $N^{\text{AM}}_{s,t,\theta,k}$ the number of active runs for school $s$ at time $\theta$ if it starts at time $h_t$, under routing scenario $k$. We can then adjust the definition of $f_{\text{cost}}(\cdot)$ to be the weighted average of the transportation cost in each scenario. Another useful extension is to allow different schedules for different days. For example, many school districts, including SFUSD, like to end classes one or two hours early on one day of the week. Such "early release" programs can allow students to participate in extra-curricular activities, and importantly, they can also provide teachers with an opportunity to host professional development activities. In Appendix B.1, we incorporate this aspect in our framework.

**Benchmarking our formulation.** Our approach makes deliberate approximations to improve tractability. It is important to evaluate their impact, as not modeling the "true" transportation costs creates optimization error as described in Section 2. We benchmark against alternatives from the literature, using large-scale simulations using Boston data to represent the true transportation costs. This allows us to make fair comparisons with Bertsimas, Delarue, and Martin (2019), who are closest to our work and use Boston data. We present the main aspects of these simulations, leaving the details to Appendix B.2. A public school transportation dataset from Boston Public Schools includes 134 schools and roughly 20,000 students. To add variety, we create 30 synthetic school district datasets from this data by sampling subsets of 25, 50, and 100 schools 10 times each. We set $T = 3$ and consider four scenarios with the allowed start times 40, 50, 60, and 70 minutes apart. In this environment, we create an optimization-based transportation simulator to make realistic routing decisions, improving on the one in Bertsimas, Delarue, and Martin (2019). Given a set of start times, we first cluster students into bus stops, then create runs of varying lengths for each school, and then schedule a bus for each run to minimize the total number of buses needed. The input of this routing engine simulator is a district (schools and students) and a start time for each school. The output is a feasible bus routing solution using as few buses as possible.

We compare our formulation for start time optimization with three approaches from Bertsimas, Delarue, and Martin (2019): (i) a "balanced" baseline which minimizes the maximum number of routes arriving at each of the three allowed start times; (ii) a "random restarts" method in which random start times are sampled for each school $N = 16$ times, then we keep the solution with the lowest transportation costs; (iii) the quadratic assignment formulation, the main approach of Bertsimas, Delarue, and Martin (2019), which seeks to maximize both the "compatibility" of runs (runs are more compatible if they can be served by the same bus with minimal deadhead) and the "spread" of start times. The last method requires tuning two key hyperparameters, for which we use a small grid with 16 parameter values. As a result, methods (ii) and (iii) require 16 transportation simulations each (to pick the best solution in the former case and optimize the hyperparameters in the latter). Results are presented in Figure 4, comparing the two baselines, the quadratic assignment formulation, and our tensor-based linear formulation. On the left panel, the quadratic formulation tends to lead to solutions with fewer buses than the balanced and random restarts baseline, as observed by Bertsimas, Delarue, and Martin 2019. How-

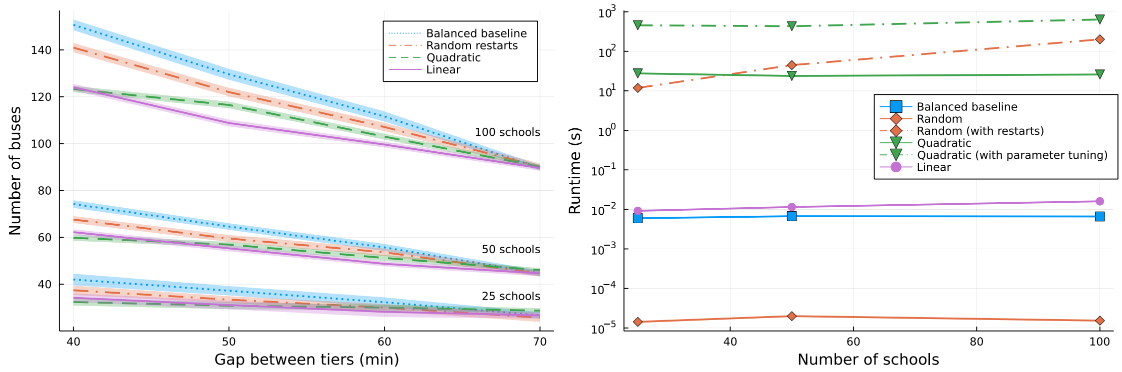


**(a)** Transportation cost **(b)** Runtime

Figure 4: Comparing performance of start time optimization algorithms.
*Note.* Our linear formulation is competitive with and sometimes beats the quadratic formulation of Bertsimas, Delarue, and Martin 2019 in terms of cost and is faster by orders of magnitude. Ribbons in the left panel indicate standard errors.

ever, our tensor-based linear formulation competes with and sometimes outperforms the quadratic formulation. As the gap between tiers grows (making connecting routes easier and the linear formulation optimal), all methods perform similarly. The right panel shows that our linear formulation has a considerable tractability edge. It is slightly slower than the balanced baseline but three orders of magnitude faster than random restarts and the quadratic formulation; parameter tuning then slows down this last method by another 1-2 orders of magnitude. Note that the quadratic assignment formulation is slow to solve, even though (like Bertsimas, Delarue, and Martin 2019), we do not solve it to optimality but instead use a local improvement heuristic where we randomly fix all but 20 start times at each iteration. We verify in the appendix that parameter tuning is indeed necessary for the quadratic method, providing improvements of up to 20%. We also observe that the solutions produced by both algorithms can be quite different despite having similar costs (evidence of flexibility for policymakers to optimize for different objectives without compromising bus costs). This significant tractability improvement over more complex formulations enables our computationally intensive interactive optimization approach, as discussed in the next section.

# 4 Interactive optimization

We can solve the computationally efficient optimization model (9) and obtain school start times almost instantly, given the objective weights $(\lambda_{\text{cost}}, \lambda_{\text{change}})$. And we can potentially add a variety of other objectives/constraints. Using the terminology of Section 2, researchers can use this model to implement the district's policy efficiently and with low optimization error. We now discuss the policymakers' side of the problem.

The most straightforward situation is when the policymakers know their objectives perfectly ex-ante, and there is no policy error. For example, if a district has the sole objective of minimizing transportation cost, then Section 3 is sufficient to address the problem. Proposition 1 states that the optimization error is the only error in the research-policy collaboration process in that case, and the benchmarking result in Section 3 indicates that it is relatively small. However, start time choice is a complex policy topic, and SFUSD did not have a clear objective ex-ante due to objective prioritization and discovery challenges.

Suppose now that the main challenge of policymakers is *objective prioritization*. That is, they know the objectives ex-ante but not how to prioritize them. For example, in Eq. (9), the policymakers do not know the values $\lambda_{\text{cost}}$, $\lambda_{\text{change}}$ but know that minimizing the transportation costs and the average change are the only two objectives they care about. In that case, a typical approach is for the researchers to compute Pareto-optimal solutions between the objectives so that policymakers can relate the abstract weights to practical outcomes and better choose them. Given that our approach is computationally fast, this is easy to achieve if the number of objectives is manageable: we can vary the weights of the various objectives ($\lambda_{\text{cost}}$, $\lambda_{\text{change}}$ in our example) and solve Eq. (9) to obtain Pareto-optimal solutions.[3] We discuss in Section 5 how this process unfolded with SFUSD. It involves the back-and-forth between researchers and policymakers modeled in Section 2: indeed, as researchers show Pareto curves and solutions to policymakers, *objective discovery* happens, and the policymakers add new objectives that were not previously considered, which requires new research and further iterations.

We facilitate this objective discovery and limit the back-and-forth using an interactive optimization approach. At a high level, we generate a set of near-optimal solutions with respect to the *known* objectives that are as diverse as possible, acknowledging that the current modeled objectives do not fully reflect all the potential true objectives of the district. We design interactive tools so policymakers can directly search this solution space and explore and compare the various solutions. This adds transparency to the collaboration process and allows policymakers to refine their preferences further (i.e., to reduce policy error) without as many back-and-forths with researchers (reducing the policy dependency on research). Specifically, suppose that the district already has a clear set of objectives and prioritization, for example minimizing $\lambda_{\text{cost}} f_{\text{cost}}(\mathbf{x}) + \lambda_{\text{change}} f_{\text{change}}(\mathbf{x})$ as in Eq. (9). In high-dimensional problems such as start time choice, the set of optimal or near-optimal solutions can be quite large, and very different solutions can have nearly identical objective values. For example, in Appendix C, Table 2 presents two sets of start times for SFUSD: the first needs 133 buses and an average change of 27.7 minutes; the second needs 135 buses

[3] Technically, the weighted-sum of objectives method can only compute the convex hull of the Pareto frontier. Other methods exist to compute nonconvex Pareto frontiers. In our context, the results are essentially the same.

and an average change of 27.0 minutes. The two schedules differ in the start times of 12 schools (9% of total). Our approach aims to showcase this flexibility to the policymakers to help them refine their preferences and reach a more informed consensus. It is rooted in Theorem 2, which demonstrates that reducing the dependency between researchers and policymakers enhances the effectiveness of their collaboration.

In what follows, Section 4.1 first introduces a general sampling approach that balances the variety of the near-optimal solutions generated (to help with objective discovery) with how sub-optimal they are (to match the known objectives). We use it to generate thousands of such solutions, leveraging the fast optimization model given in Section 3. These solutions are loaded into an interactive start time choice interface, described in Section 4.2, that aims to facilitate the objective discovery process and, more generally, increase the transparency of the optimization model. Then, Section 5 discusses how SFUSD used this approach to support the asynchronous work between researchers and policymakers and the specific policy improvements it unlocked.

### 4.1 Sampling the near-optimal solution space

We assume the policymakers' current policy preferences are modeled by an objective function $f^*$. In the previous example, we have $\lambda_{\text{cost}} f_{\text{cost}}(\mathbf{x}) + \lambda_{\text{change}} f_{\text{change}}(\mathbf{x})$ for a choice of weights. We denote the optimal objective value given these weights as $c^* = \min_{\mathbf{x}\in\mathcal{X}} f^*(\mathbf{x})$. In our illustrative model, this objective maps to the current iteration of the policy decision. We aim to allow the district to explore the near-optimal solution space, understand the variety of existing solutions, and discover new objectives. The first step is to sample and compute diverse near-optimal solutions. We achieve this by using a random matrix $P$, in which each element $P_{s,t}$ represents a preference score for school $s$ to start at time $t$ and is sampled uniformly and independently between 0 and 1. We then solve the following optimization problem for each realization of $P$:

$$\begin{aligned} \max_{\boldsymbol{x}\in\mathcal{X}} \quad & \sum_{s=1}^{S}\sum_{t=1}^{T} P_{s,t} x_{s,t} \\ \text{s.t.} \quad & f^*(\mathbf{x}) \leq (1+\gamma)c^*, \end{aligned}$$

where $\gamma > 0$ is a budget multiplier that controls the maximum deterioration in the original objective $c^*$ that we are willing to accept. The random preference matrix helps to generate solutions that are as varied as possible, while $\gamma$ controls the tradeoff between solution variety and suboptimality. We repeat this sample-then-optimize process $N$ times and denote the optimal solution as $\mathbf{x}^i$ for each sample $i = 1, \ldots, N$. $H_{i,s} = \sum_t x^i_{s,t} h_t$ is the start time of school $s$ in the sample solution $i$. a A key challenge is to choose a value of $\gamma$ that carefully balances solution quality and variety: generating $N$ solutions is not very helpful if they are all identical, and we want solutions that place individual schools at as many different start times as possible. To measure the variety of the solutions we obtain,

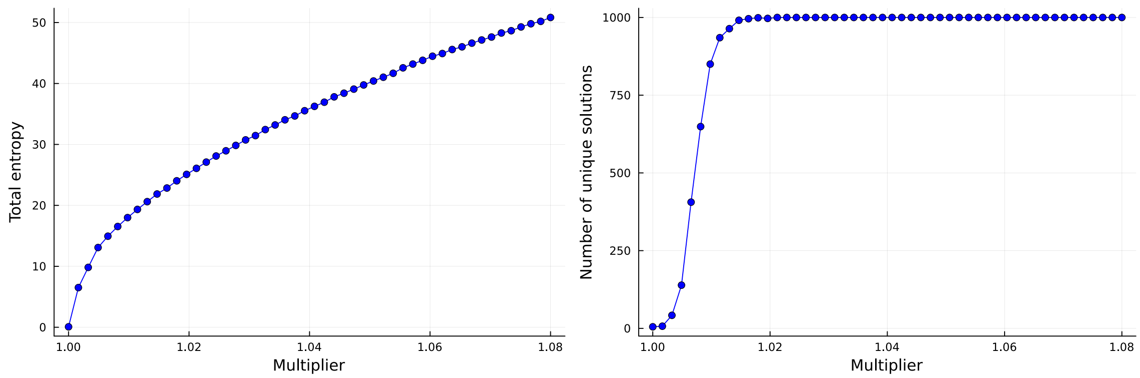


**(a)** Total entropy $E$ **(b)** Number of unique solutions

Figure 5: Effect of the budget multiplier on solution diversity for a sample of $N = 1,000$ solutions.

we let $p(s,t) = \frac{1}{N}\sum_i \mathbf{1}_{H_{i,s}=h_t}$. Here, $p(s,\cdot)$ is the probability mass function of the empirical distribution of the start time of school $s$ across our sampled near-optimal solutions. The Shannon entropy of this distribution is $e_s = -\sum_t p(s,t)\ln p(s,t)$: it is equal to 0 if school $s$ is fixed at the same time in all $N$ samples, and $e_s = \ln 3$ if school $s$ is equally likely to be in any of the three start times. We sum the entropy $e_s$ of each school, such that $E = \sum_{s=1}^{S} e_s$ is our measure of the diversity of sampled solutions. It has a simple interpretation: as $\ln 3 \approx 1$, we can roughly interpret $E$ as the number of schools that are not fixed in our set of near-optimal solutions. Indeed, if all solutions are identical, we have $E = 0$, while $E = S\ln 3$ when all schools are equally likely to get each start time across our solutions.

We use this metric to select the budget multiplier $\gamma$. Because the optimization model from Section 3 is tractable, we can sample thousands of solutions for many different values of $\gamma$. As $\gamma$ increases, we move further away from the optimum of $f^*$, and solution diversity $E$ increases. Figure 5 showcases the relationship between $\gamma$ and $E$ for a particular problem we faced with SFUSD (discussed in Section 5). This figure is representative of our experience optimizing start times in large districts. It shows significant variety in near-optimal solutions: a $\gamma = 2\%$ optimality gap is enough to ensure that all 1,000 sampled solutions are distinct, and the solutions get increasingly diverse as we increase $\gamma$. We select the smallest value of $\gamma$ that creates a rich enough solution space. Policymakers can explore these solutions to discover what is achievable, enable further discussions, and refine their policy. However, this process requires an accessible way to navigate multiple solutions.

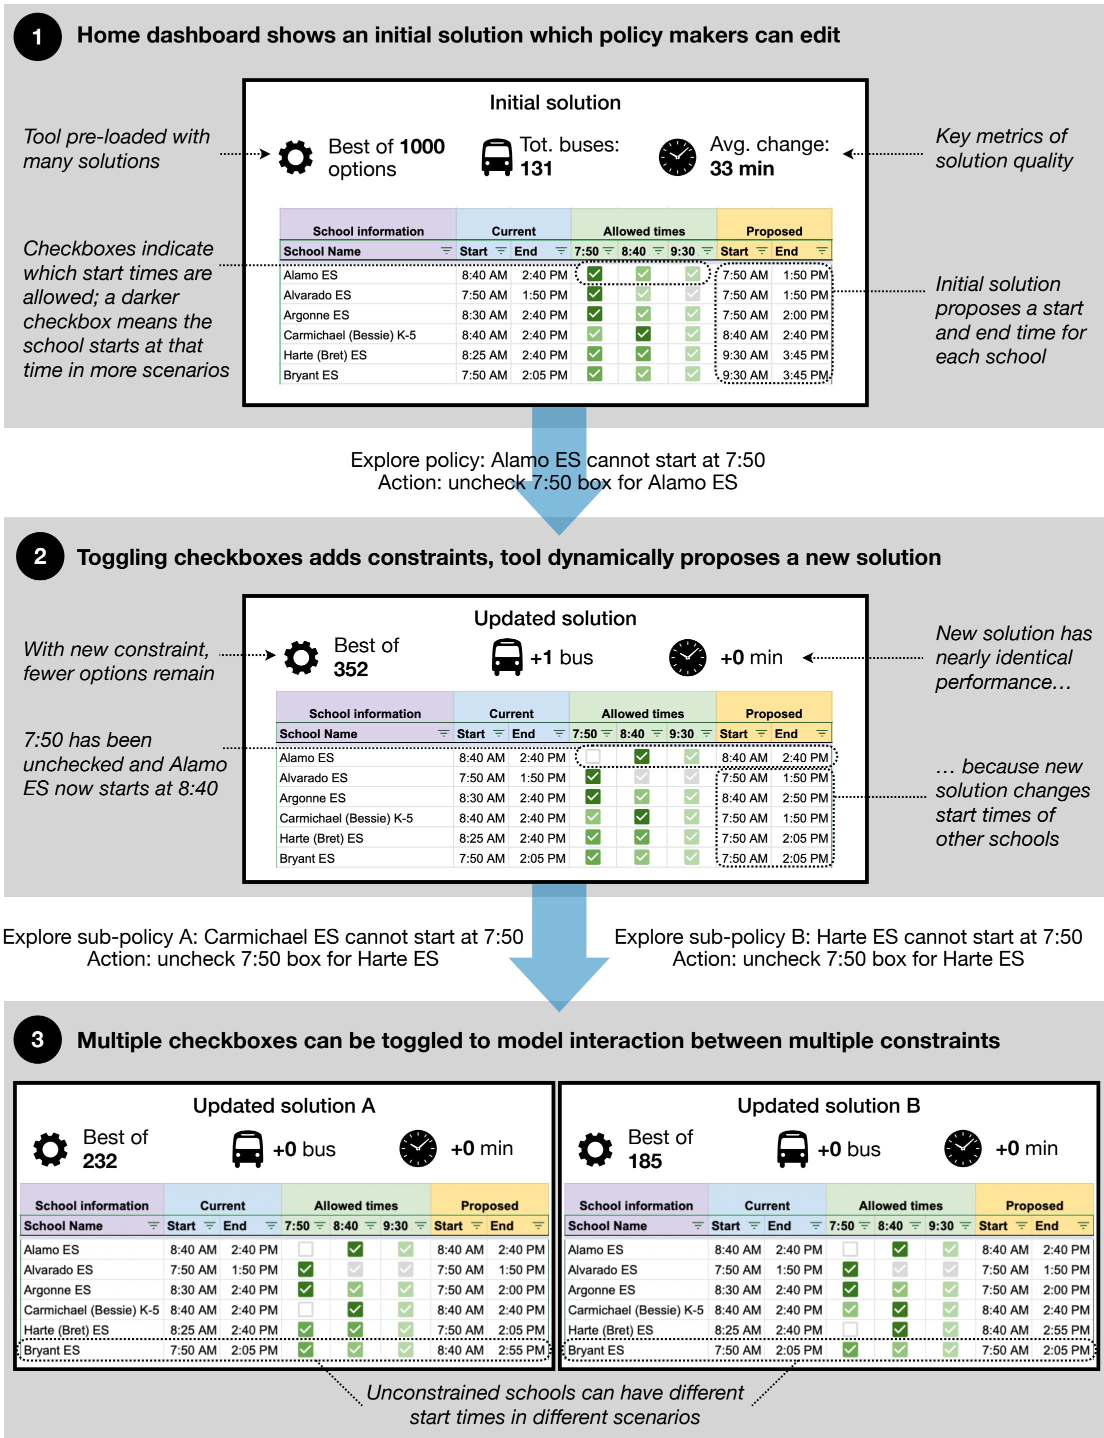


Figure 6: Diagram of our interactive optimization tool.

### 4.2 An interactive optimization interface

In order to allow for solution space exploration at SFUSD, we designed an interactive start time choice interface, illustrated in Figure 6. This tool is pre-loaded with the previously described $N = 1000$ diverse solutions. As displaying all solutions is impractical, we only show the best-sampled solution $j^* \in \{1, \cdots, N\}$ according to the current district objective $f^*$. That is, $j^* = \arg\min_{i \leq N} f^*(\mathbf{x}^i)$. The objective values for each solution $j$ are precomputed, making finding $j^*$ trivial. Box 1 of Figure 6 shows how we display the solution: we list the proposed start and end time of all schools, together with a summary of the solution's objectives (here $f_{\text{cost}}$ and $f_{\text{change}}$). We also include an interactive map to help stakeholders quickly identify start times at schools they care about. Notice the green checkboxes (one for each start time) for each school in the "allowed times" column of the figure. Checkbox shading for school $s$ represents its empirical distribution $p(s, \cdot)$, quickly showing the policymakers which schools are more or less flexible. Then, policymakers can uncheck one or more boxes to disallow certain times for certain schools. Unchecking the checkbox for school $s_0$ at start time $h_{t_0}$ imposes the constraint $x_{s_0,t_0} = 0$, and updates the displayed best solution to $j^{*\prime}$, with

$$j^{*\prime} = \arg \min_{i \leq N, \text{ s.t. } x^i_{s_0,t_0}=0} f^*(\mathbf{x}^i),$$

if it is feasible. More generally, pre-loaded solutions that break the checkbox constraints are marked as infeasible and henceforth disregarded. Boxes 2 and 3 of Figure 6 illustrate how policymakers use this functionality and note how the solution and empirical distribution in the green checkboxes are updated dynamically. This tool is entirely encoded as a Google spreadsheet, as it is designed to be highly portable and accessible to any stakeholder or policymaker with a computer and an internet browser.

Policymakers are, therefore, able to manually explore the near-optimal solution space. In particular, changing the time of one school with the tool will typically trigger the change of many other schools to keep the solution near-optimal. For example, moving a school to a later time is likely to force at least one other school to move earlier. This design choice addresses a fundamental policy hurdle. Many stakeholders' interest in any start-time solution is limited to a few schools. For example, individual families or staff members may have strong opinions about start and end times at the schools that employ them or educate their children, while remaining indifferent to changes at schools with which they are not affiliated. This single-school "tunnel vision" can make it difficult for planners to align on a global policy. If one school's stakeholders really want to change its start time, the interactive optimization tool can immediately show the ripple effects on potentially many other schools, enabling a productive conversation. Because of the complexity of school transportation, central planners (let alone individual stakeholders) have only a limited understanding of these ripple effects, which is why the tool's immediate feedback is helpful.

There are limits to this approach. If the policymakers make too many changes, we may run out of precomputed solutions that satisfy their modifications, as shown by the reduced solution count in Figure 6 as constraints are added. If that is the case, policymakers can simply meet with the researchers again and explain what they were trying to achieve. This is then translated into an updated formulation with additional objectives and constraints and new objective prioritization, and the researchers can run the solver to sample $N$ new near-optimal solutions, which update the interactive interface, repeating the process. This update takes time, and it corresponds to the "lower research speed, lower policy dependency on research" motivated by Section 2. Another limitation is that the district may sometimes want to change a school's time without triggering any other changes and simply understand the resulting cost increases. For this application, we created a *sensitivity analysis tool*, which affords policymakers more control and transparency over the optimization process at the expense of optimality. This tool, detailed in Appendix D and Figure 14, is also built to provide as much insight as possible into the consequences of district choices on its transportation system. Given the motivating intuition from Section 2 and the guiding principles of the interactive optimization interface, we believe that our approach could be generalized to other complex policy-and-optimization problems.

The tools introduced in this section only have value insofar as they enhance the policy-making process. We next describe how SFUSD used our approach to update start times and provide evidence of its benefits.

# 5 Changing school start times in San Francisco

This section details how our interactive framework enabled the successful change of SFUSD's start times. We highlight the interplay between research and policy, showcasing the process of objective discovery. We also analyze a large-scale post-implementation survey of SFUSD's staff and parents to provide evidence of the benefits of our approach.

## 5.1 Background: Winds of change in the Bay.

Our collaboration with SFUSD began in December 2020, when the district — the seventh-largest in California, enrolling over 50,000 students across more than 120 schools — planned a comprehensive overhaul of its school schedules (SFUSD 2021). In the past, individual schools had been given significant scheduling autonomy, leading to 18 different start times in the 2019-2020 school year. Afternoon schedules were even more diverse, with a majority of schools offering "early release" on at least one day of the week. The lack of consistency across schools made school bus routing more complex. SFUSD was only responsible for transporting special education students (less than 10% of total enrollment); however, these students were spread across most schools in the district, meaning no school was immune from transportation considerations. Additionally, like most urban school districts in California, SFUSD did not yet comply with State Senate Bill 328 (SB-328), which mandated

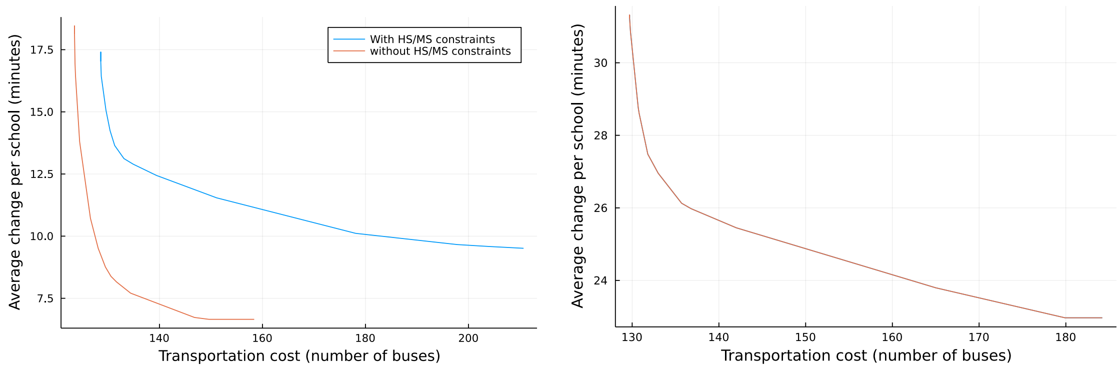


**(a)** First Pareto curves shared with SFUSD **(b)** Pareto curve with SFUSD's modifications

Figure 7: Visualizing tradeoffs between transportation costs and average change
*Note.* The left plot is computed with and without legal constraints (SB-328), and the right plot includes them.

later start times for middle and high schools starting in the 2022-2023 school year.

Due to the COVID-19 pandemic, SFUSD had moved to remote learning starting in March 2020. The 2020-2021 school year also started online, with some schools gradually resuming in-person activities in the spring of 2021. While orchestrating the reopening of every school in the district after an unprecedented district-wide shutdown was no simple task, it also afforded SFUSD a chance to reset its operations, including school start and end times. Indeed, COVID disruptions had reduced the ordinary pull of the status quo and created an unprecedented opportunity for districtwide scheduling changes.

## 5.2 Research and policy collaboration with interactive optimization

From December 2020 to April 2021, the research team engaged in weekly meetings with key SFUSD stakeholders to exchange ideas and gather feedback. When we first started to meet, the district was creating a new set of start times, and transportation experts were putting in weeks of manual work to evaluate its impact on transportation. As we presented them with automated ways to estimate transportation costs and optimize start times, the district leaders leveraged this new freedom. They moved their focus to understanding the tradeoff between the two objectives of maximizing transportation savings and minimizing change, which is why these objectives are featured prominently in Section 3. Assessing this tradeoff was further complicated because high schools and middle schools were forced to move to comply with SB-328. Figure 7a shows one of the first Pareto curves we shared with SFUSD to help them quantify this tradeoff. The relationship between the average change in school start times (in minutes) and SFUSD's total transportation cost (in buses), both with and without compliance with SB-328. We observe that it is possible to reduce bus

Figure 8: Optimizing transportation costs and start time change

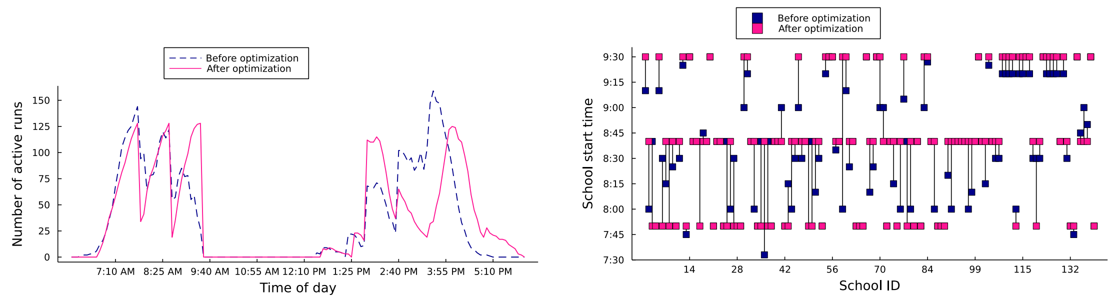


**(a)** Number of active runs over time, before and after minimizing cost and change

**(b)** School start time schedules, before and after minimizing cost and change

costs without a drastic change in school start times; moreover, imposing legal constraints on high school and middle school start times can carry either a huge cost penalty (50 buses or more) or a moderate change penalty (5-10 more minutes of average change across the district). The bottom right point of the unconstrained curve corresponds to the nearest solution to the status quo, which was infeasible since some schools did not start at one of the three allowed start times. These curves are obtained with the optimization approach from Section 3, and Figure 8 illustrates one particular solution. Figure 8a shows the number of active buses throughout the day, defined as $(N^{\mathrm{AM}}_{s,t,\theta} + N^{\mathrm{PM}}_{s,t,\theta})x_{s,t}$ for each time $\theta$. In the status quo (blue dashed line), SFUSD has unbalanced peaks in the morning and afternoon, leading to high transportation costs (159 buses). Recall that $f_{\text{cost}}$ from Equation (11) minimizes the maximum number of buses used in the day. This is why, even with the additional objective of minimizing changes, the number of active runs is perfectly balanced into three tiers in the morning, reducing transportation costs (128 buses). Figure 8b also shows how the model reduces change; the blue squares represent start times before optimization and the pink squares after optimization. The black lines represent the relative change in start time for a given school: to minimize change, many (but not all) schools move to the nearest allowed start time, and only one start time changes by more than 50 minutes.

As described in Section 2, the interplay between research and policy required many iterations. For example, after considering the results in Figure 7a, SFUSD assessed that they were willing to accept more change to secure some other objectives. They chose to add some constraints to the model, such as moving all high schools to the same start time out of fairness concerns, as well as some school-specific constraints. In Figure 7b, we observe the tradeoff between transportation cost and average change per school with these new constraints. While the cost objective has roughly the same range of values, the change objective has roughly doubled from around 12 minutes to around 25 minutes per school. After several iterations, the policymakers converged on targeting 130 buses and 30 minutes of average change. We could have stopped at this point because iterating further with more objectives would require many more weeks or months of exchange. Instead, we designed the interactive optimization approach from Section 4 to enable SFUSD to understand which solutions could best align with other priorities beyond cost, change, and the aforementioned operational constraints. SFUSD leaders found the tool simple and intuitive. Without our help, they made many manual changes to the solution that were highly nontrivial. Our only involvement was to update the tool with the district's updated choices and generate new solutions to help policymakers explore further. The district also used the "manual optimization" sensitivity analysis tool from Appendix D to make local changes that were impossible with the other optimization model. Overall, these few weeks with the interactive tools were the most productive on the policy side. Policymakers quickly iterated without us and changed the time of 25 schools such that:

- the schedules of nearby elementary schools were aligned to allow for joint professional development activities for teachers and staff ("common planning time"),
- the weekly instructional time at each school was tweaked to create more homogenous end-time patterns, improving fairness and transportation,
- the last-minute requests of various schools and stakeholders were carefully considered and sometimes incorporated,
- the change was more fair: the number of schools shifting by more than 50 minutes was reduced from 7 to 1, while 9 schools with less change shifted more.

Note that these examples of changes are either completely novel objectives (the common planning time) or modeling variations on existing objectives (the amount of change). This second point illustrates the need for interactivity; even for a relatively clearly defined objective like the level of change, different metrics (e.g., average change vs. max change) can lead to drastically different solutions, and the researchers cannot specify the perfect metric (if one even exists) ex-ante. The interactive approach goes a long way in capturing objectives not fully represented by the optimization program.

By April 2021, the school district had new start times for their 133 public schools. Using a policy-research collaboration and interactive optimization, the district could assign each school to one of three start times (7:50, 8:40, 9:30). The new start times complied with SB-328, one year ahead of schedule. In addition, the solution aligned afternoon end times across many nearby elementary schools to allow for teacher professional development. The new start times granted the district nearly $5.5 million in annual transportation savings.

### 5.3 Post-implementation survey

A key remaining question is to evaluate the success of this approach in a principled way. Time has provided some evidence of this, as the new start times remain in use two years after the change with minimal modifications. The year after implementation, we also used a survey to understand families' and staff's satisfaction. Before the start time change in February 2021, we had designed a survey to gather feedback from parents and staff and incorporate it into the selection process. However, SFUSD chose not to distribute it, largely due to concerns about survey fatigue during COVID. Consequently, preferences from parents and staff regarding start times were not directly measured before the change. Indeed, the lack of knowledge about detailed stakeholder preferences was one of the motivations behind our interactive optimization approach in the first place. Our goal was to empower policymakers, who were directly in contact with stakeholders and knew their preferences, to make decisions that would best satisfy them. We therefore sought to understand to what extent this goal was achieved in practice.

In April 2022, we designed a survey for SFUSD to obtain feedback from parents and staff across schools affected by the schedule change in Fall 2021. This survey was distributed to approximately 24,000 families and 3,551 staff from 70 schools. The survey received about 6,400 responses from parents, achieving a 26.7% response rate, and 1,522 responses from staff, achieving a 43.7% response rate.[4] The results must be interpreted carefully, as the survey was conducted *after* the schedule change, which means the reported preferences over start times may be subject to bias, including post-hoc rationalization (favoring the implemented school schedule) and responder selection bias. We acknowledge these potential issues and try to provide a balanced analysis of the survey results. Appendix E.2 in the appendix includes a sample and visualization of the survey questions.

The first question of the parent survey asked respondents to indicate their satisfaction level with their school's new start time on a five-point scale from "very unsatisfied" to "very satisfied". Results are shown in Figure 9. Overall, parent respondents are positive about the change. Figure 9a shows a slight majority of parents are at least somewhat satisfied while 37% are very satisfied. However, about 23% of respondents were "very unsatisfied," which illustrates how hard these changes are on families, and is in line with the difficulties experienced by other districts in changing start times. To better understand the unsatisfied parents, the survey's second question asked respondents to rank the three possible start time options in order of personal preference. Figure 9b correlates these responses to the first question (about satisfaction with the outcome). The x-axis represents a respondent's ranking of the school start time under the new schedule: for example, if a parent has an 8:40 AM start time under the new schedule, and she ranks 8:40 AM as her top choice in the survey, then her response will be counted as one instance under the "First" group. The y-axis then counts the number of instances in each rank group, and each group is further broken down by the satisfaction level shown by the different color bars.

Surprisingly, "very unsatisfied" respondents were almost as likely to have received their favorite start time as their least favorite start time. In contrast, "very satisfied" respondents

[4] After dropping the responses with incomplete answers, the total number of effective responses from parents is 5,773; the total number of effective responses from the staff is 1,314.

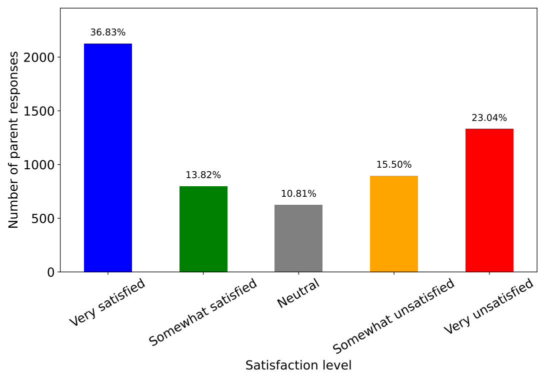


**(a)** Distribution of parent satisfaction

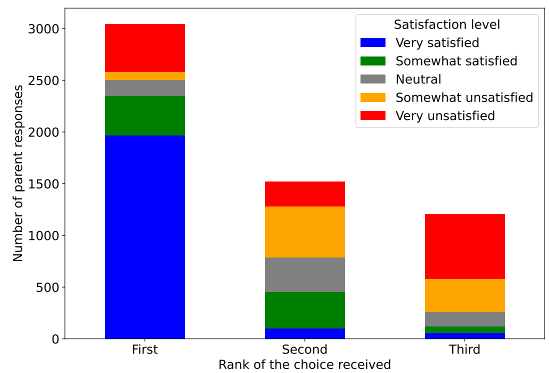


**(b)** Parent satisfaction by rank of the choice received

Figure 9: Parent survey results and satisfaction levels

*Note.* Survey results from 5,773 parents. The left panel shows responses to the first question, which asked parents to rate their satisfaction level with their new start time on a five-point scale. The right panel crosses the responses to this first question with a second question asking parents to rank the three possible start times that they *could* have received; for instance, the left-most bar counts all responses from parents who ranked their new start time as their favorite option in absolute terms.

overwhelmingly received their first choice. We propose two possible explanations for this discrepancy. The first is that "very unsatisfied" respondents would have preferred another start time option besides the three selected by SFUSD — for instance, starting school at 8:10 rather than 7:50 or 8:40. The second is that respondents might be using the survey to voice dissatisfaction about broader issues than start time selection — always a possibility in such surveys. For example, parents with two children in different schools who received their preferred start time for one child but not the other might indicate dissatisfaction in both of their survey responses. Either way, analyzing dissatisfied respondents reflects the ultimate challenge of policy-making: satisfying all members of a diverse population is inherently difficult, and even the best possible aggregate solution may still disappoint some individuals. The staff survey results are shared in Appendix E.1 and are very similar to the parent survey: a plurality (31.89%) of staff members report being "very satisfied" about the change.

## 5.4 Evaluating the benefits of interactive optimization

Table 1: Cost comparison between two sets of schedules

| | Transportation cost (buses) | Average change (minutes) |
|---|---|---|
| Final schedule | 134 | 32.82 |
| Min-cost schedule | 131 | 33.63 |

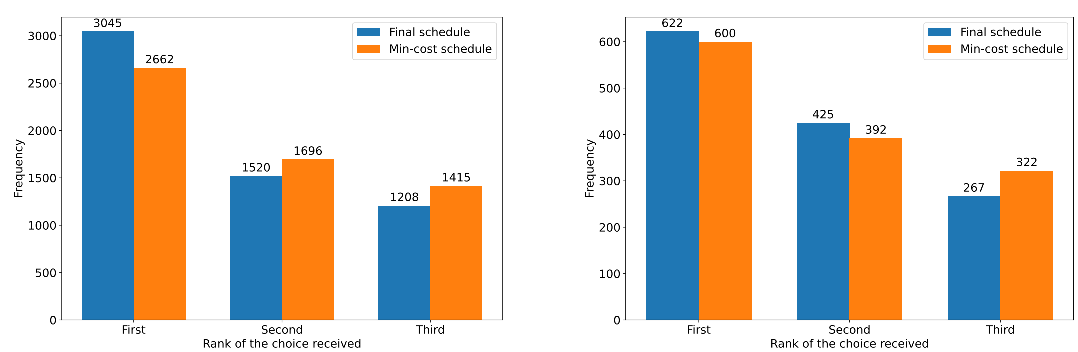


**(a)** Parents **(b)** Staff

Figure 10: Comparing parent rankings of the final and minimum-cost schedules
*Note.* Each histogram shows the number of respondents receiving their first, second, and third choices under each proposed schedule. The corresponding cost for each set of schedules is shown in Table 1.

The choice ranking data behind Figure 9b also allow us to run a counterfactual analysis.

Recall that before SFUSD started to use our interactive tools, they had already chosen a tradeoff between transportation costs and the amount of change after several iterations of Pareto curve analysis. Without the interactive approach, the district might have stuck with these objectives and chosen the times Gurobi (our optimization solver) recommended instead of the many modifications they made with the interactive software (recall that the times of 25 schools were changed). In Table 1, we compare the implemented schedule ("final schedule", with interactive optimization) and the counterfactual output of Gurobi before the manual changes ("Min-cost schedule"). The table shows they have similar transportation costs and average change despite the many updates; this is a feature of the interactive approach, which only considers near-optimal solutions. Figure 10 uses the ranking data from the survey response to evaluate how the counterfactual schedule (without interactive optimization) would have affected families and staff. This comparison shows that 14% more parents received their first choice than in the counterfactual, as well as a 17% decrease in the number of staff members getting their last choice. While statistically significant, these results suffer from the many limitations of survey data: the survey was conducted after the change, and there may be selection bias in our survey responses.

Ultimately, this paper presents diverse evidence in favor of interactive optimization. This evidence includes the survey results; the many community-driven changes implemented by SFUSD and listed in Section 5.2; the project's status as the first successful optimization-driven school start time change in the US and its continued implementation at the time of writing; and the structural analysis from Section 2. Another piece of evidence is the comparison with the approach of Bertsimas, Delarue, Eger, et al. 2020, who collaborated with Boston Public Schools (a district of comparable size) on a similar project in 2017. They also designed a multi-objective bell time optimization algorithm, but the solution recommended to the district received significant pushback and was ultimately shelved. According to Goodman 2019, implementation failed because "there was almost no engagement with the model itself, insufficient transparency about the algorithm's tradeoffs, and no opportunity to adjust it." While comparing the projects head-to-head is a risky exercise in small-sample statistics, it is fair to say that the successful implementation in SFUSD depended not only on modeling accuracy but, more importantly, on policymakers' ability to make adjustments as needed. Human input not only improved the solution by avoiding potential blind spots in the optimization model; it also built trust and ownership in the policymakers who bore the responsibility of the decision. In our multi-step approach, our algorithms gradually ceded control of the solution to the human experts. We are optimistic that such a framework may have merit in other complex problems at the intersection of operations and policy.

Reflecting on the outcomes of our partnership, a key policymaker from SFUSD shared:

> "This research has been a way for us to get really advanced thinking and solutions to help us address complex problems. With the researchers' help, SFUSD

was able to make school schedule changes that created more opportunities for aligned planning at school sites for strong teaching and learning, and opened the door for $5 million worth of transportation efficiencies. That's more money for students."

## References


Ahuja, Ravindra K, Krishna C Jha, and Jian Liu (2007). "Solving real-life railroad blocking problems". In: *Interfaces* 37.5, pp. 404–419.

American Medical Association (2016). *Insufficient Sleep in Adolescents*.

Bandaru, Sunith, Amos HC Ng, and Kalyanmoy Deb (2017). "Data mining methods for knowledge discovery in multi-objective optimization: Part A-Survey". In: *Expert Systems with Applications* 70, pp. 139–159.

Banerjee, Dipayan and Karen Smilowitz (2019). "Incorporating equity into the school bus scheduling problem". In: *Transportation Research Part E: Logistics and Transportation Review* 131.October, pp. 228–246. ISSN: 13665545. arXiv: 1811.11322.

Barbati, Maria, Salvatore Corrente, and Salvatore Greco (2024). "Multiobjective Combinatorial Optimization with Interactive Evolutionary Algorithms: the case of facility location problems". In: *EURO Journal on Decision Processes*, p. 100047.

Bertsimas, Dimitris, Arthur Delarue, William Eger, et al. (Jan. 2020). "Bus Routing Optimization Helps Boston Public Schools Design Better Policies". In: *INFORMS Journal on Applied Analytics* 50.1, pp. 37–49. ISSN: 0092-2102.

Bertsimas, Dimitris, Arthur Delarue, and Sebastien Martin (2019). "Optimizing schools' start time and bus routes". In: *Proceedings of the National Academy of Sciences of the United States of America* 116.13, pp. 5943–5948.

Bolton, Gary E and Elena Katok (2018). "Cry wolf or equivocate? Credible forecast guidance in a cost-loss game". In: *Management Science* 64.3, pp. 1440–1457.

Brill, E Downey, Shoou-Yuh Chang, and Lewis D Hopkins (1982). "Modeling to generate alternatives: The HSJ approach and an illustration using a problem in land use planning". In: *Management Science* 28.3, pp. 221–235.

Buell, Ryan W, Tami Kim, and Chia-Jung Tsay (2017). "Creating reciprocal value through operational transparency". In: *Management Science* 63.6, pp. 1673–1695.

Caro, Felipe et al. (2004). "School redistricting: Embedding GIS tools with integer programming". In: *Journal of the Operational Research Society* 55.8, pp. 836–849.

Carrell, Scott E., Teny Maghakian, and James E. West (2011). "A's from Zzzz's? The causal effect of school start time on the academic achievement of adolescents". In: *American Economic Journal: Economic Policy* 3.3, pp. 62–81. ISSN: 19457731.

Chang, Shoou-Yuh, E Downey Brill Jr, and Lewis D Hopkins (1982). "Use of mathematical models to generate alternative solutions to water resources planning problems". In: *Water Resources Research* 18.1, pp. 58–64.

Chu, Amanda, Pinar Keskinocak, and Monica C Villarreal (2020). "Empowering Denver Public Schools to Optimize School Bus Operations". In: *INFORMS Journal on Applied Analytics* 50.5.

Crowley, Stephanie J. et al. (Aug. 2018). "An update on adolescent sleep: New evidence informing the perfect storm model". In: *Journal of Adolescence* 67.3, pp. 55–65. ISSN: 01401971.

Danna, Emilie et al. (2007). "Generating multiple solutions for mixed integer programming problems". In: *Integer Programming and Combinatorial Optimization. IPCO 2007. Lecture Notes in Computer Science*. Ed. by Matteo Fischetti and David P. Williamsom. Vol. 4513 LNCS, pp. 280–294. ISBN: 9783540727910.

Danner, Fred and Barbara Phillips (2008). "Adolescent sleep, school start times, and teen motor vehicle crashes". In: *Journal of Clinical Sleep Medicine* 4.6, pp. 533–535.

Davis, Andrew M et al. (2022). "The Best of Both Worlds: Machine Learning and Behavioral Science in Operations Management". In: *Available at SSRN 4258273*.

Desrosiers, Jacques et al. (1986). "TRANSCOL: a multi-period school bus routing and scheduling system". In: *TIMS Studies in the Management Sciences* 22, pp. 47–71.

Dietvorst, Berkeley J, Joseph P Simmons, and Cade Massey (2015). "Algorithm aversion: people erroneously avoid algorithms after seeing them err." In: *Journal of Experimental Psychology: General* 144.1, p. 114.

– (2018). "Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them". In: *Management Science* 64.3, pp. 1155–1170.

Ellegood, William A. et al. (2020). "School bus routing problem: Contemporary trends and research directions". In: *Omega (United Kingdom)* 95.xxxx. ISSN: 03050483.

Furnon, Vincent and Laurent Perron (Mar. 7, 2024). *OR-Tools Routing Library*. Version v9.9. Google. URL: `https://developers.google.com/optimization/routing/`.

Geoffrion, Arthur M, James S Dyer, and A Feinberg (1972). "An interactive approach for multi-criterion optimization, with an application to the operation of an academic department". In: *Management science* 19.4-part-1, pp. 357–368.

Gershenfeld, Gabriel (2015). "Conjoint analysis for ticket offerings at the Cleveland Indians". In: *Interfaces* 45.2, pp. 166–174.

Goldin, Andrea P. et al. (2020). "Interplay of chronotype and school timing predicts school performance". In: *Nature Human Behaviour* 4.4, pp. 387–396. ISSN: 23973374.

Goodman, Ellen P (2019). "The challenge of equitable algorithmic change". In: *The Regulatory Review, Feb.*

Greistorfer, Peter et al. (2008). "Experiments concerning sequential versus simultaneous maximization of objective function and distance". In: *Journal of Heuristics* 14.6, pp. 613–625. ISSN: 13811231.

Meltzer, Lisa J et al. (2021). "Changing school start times: impact on sleep in primary and secondary school students". In: *Sleep* April, pp. 1–14. ISSN: 0161-8105.

Miettinen, Kaisa, Francisco Ruiz, and Andrzej P Wierzbicki (2008). "Introduction to multiobjective optimization: interactive approaches". In: *Multiobjective optimization: interactive and evolutionary approaches*. Springer, pp. 27–57.

Olavson, Thomas and Chris Fry (2008). "Spreadsheet decision-support tools: lessons learned at Hewlett-Packard". In: *Interfaces* 38.4, pp. 300–310.

Owens, Judith et al. (2014). "School Start Time Change: An In-Depth Examination of School Districts in the United States". In: *Mind, Brain, and Education* 8.4, pp. 182–213. ISSN: 1751228X.

Park, Junhyuk and Byung In Kim (2010). “The school bus routing problem: A review”. In: *European Journal of Operational Research* 202.2, pp. 311–319. ISSN: 03772217.

Serra, Thiago and J. N. Hooker (2020). *Compact representation of near-optimal integer programming solutions*. Vol. 182. 1-2. Springer Berlin Heidelberg, pp. 199–232. ISBN: 1010701901390.

SFUSD (2021). *Facts about SFUSD at a glance*. URL: `https://www.sfusd.edu/about-sfusd/facts-about-sfusd-glance`.

Swersey, Arthur J. and Wilson Ballard (1984). “Scheduling School Buses.” In: *Management Science* 30.7, pp. 844–853. ISSN: 00251909.

Teghem, Jacques, Daniel Tuyttens, and Ekunda L Ulungu (2000). “An interactive heuristic method for multi-objective combinatorial optimization”. In: *Computers & Operations Research* 27.7-8, pp. 621–634.

Trapp, Andrew C. and Renata A. Konrad (2015). “Finding diverse optima and near-optima to binary integer programs”. In: *IIE Transactions (Institute of Industrial Engineers)* 47.11, pp. 1300–1312. ISSN: 15458830.

Voll, Philip et al. (2015). “The optimum is not enough: A near-optimal solution paradigm for energy systems synthesis”. In: *Energy* 82, pp. 446–456. ISSN: 03605442.

Widome, Rachel et al. (2020). “Association of Delaying School Start Time with Sleep Duration, Timing, and Quality among Adolescents”. In: *JAMA Pediatrics* 174.7, pp. 697–704. ISSN: 21686211.

Xin, Rui et al. (2022). “Exploring the whole rashomon set of sparse decision trees”. In: *Advances in neural information processing systems* 35, pp. 14071–14084.

Yang, Jian-Bo (1999). “Gradient projection and local region search for multiobjective optimisation”. In: *European journal of operational research* 112.2, pp. 432–459.

Zechman, Emily M, Marcio H Giacomoni, and M Ehsan Shafiee (2013). “An evolutionary algorithm approach to generate distinct sets of non-dominated solutions for wicked problems”. In: *Engineering Applications of Artificial Intelligence* 26.5-6, pp. 1442–1457.

Zeng, Liwei, Sunil Chopra, and Karen Smilowitz (2022). “A bounded formulation for the school bus scheduling problem”. In: *Transportation Science*.

Zionts, Stanley and Jyrki Wallenius (1983). “An interactive multiple objective linear programming method for a class of underlying nonlinear utility functions”. In: *Management science* 29.5, pp. 519–529.

# Appendices

## A Supporting content for Section 2

*Proof.* Proof of Proposition 1. We can add and subtract appropriate terms to make the policy and optimization error appear in the definition of the total error:

$$
\begin{aligned}
\varepsilon_{\text{tot}}(\mathbf{x}) &= c(\mathbf{x}, \boldsymbol{\lambda}^{\text{true}}) - c(\mathbf{x}^*(\boldsymbol{\lambda}^{\text{true}}), \boldsymbol{\lambda}^{\text{true}}) \\
&= c(\mathbf{x}, \boldsymbol{\lambda}) - c(\mathbf{x}, \boldsymbol{\lambda}) + c(\mathbf{x}^*(\boldsymbol{\lambda}), \boldsymbol{\lambda}) - c(\mathbf{x}^*(\boldsymbol{\lambda}), \boldsymbol{\lambda}) + c(\mathbf{x}^*(\boldsymbol{\lambda}), \boldsymbol{\lambda}^{\text{true}}) - c(\mathbf{x}^*(\boldsymbol{\lambda}), \boldsymbol{\lambda}^{\text{true}}) \\
&+ c(\mathbf{x}, \boldsymbol{\lambda}^{\text{true}}) - c(\mathbf{x}^*(\boldsymbol{\lambda}^{\text{true}}), \boldsymbol{\lambda}^{\text{true}}) \\
&= \varepsilon_{\text{tot}}(\mathbf{x}, \boldsymbol{\lambda}) + \varepsilon_{\text{pol}}(\boldsymbol{\lambda}) + [c(\mathbf{x}, \boldsymbol{\lambda}^{\text{true}}) - c(\mathbf{x}, \boldsymbol{\lambda})] - [c(\mathbf{x}^*(\boldsymbol{\lambda}), \boldsymbol{\lambda}^{\text{true}}) - c(\mathbf{x}^*(\boldsymbol{\lambda}), \boldsymbol{\lambda})].
\end{aligned}
$$

We define the third term as $\delta(\mathbf{x}, \boldsymbol{\lambda}) = [c(\mathbf{x}, \boldsymbol{\lambda}^{\text{true}}) - c(\mathbf{x}, \boldsymbol{\lambda})]$. As for the fourth term, $[c(\mathbf{x}^*(\boldsymbol{\lambda}), \boldsymbol{\lambda}^{\text{true}}) - c(\mathbf{x}^*(\boldsymbol{\lambda}), \boldsymbol{\lambda})]$, we notice that we can always add an appropriate constant to each cost function $c(\cdot, \boldsymbol{\lambda})$ such that for every $\boldsymbol{\lambda}$, $c(\mathbf{x}^*(\boldsymbol{\lambda}), \boldsymbol{\lambda}) = c(\mathbf{x}^*(\boldsymbol{\lambda}), \boldsymbol{\lambda}^{\text{true}})$. □

*Proof.* Proof of Lemma 1. Recall the model dynamics:

$$
\begin{aligned}
\varepsilon_{\text{pol}}^{n+1} &= (1-\pi)\varepsilon_{\text{pol}}^{n} + \pi\gamma\varepsilon_{\text{opt}}^{n} \\
\varepsilon_{\text{opt}}^{n+1} &= \varepsilon_{\text{pol}}^{n} - \varepsilon_{\text{pol}}^{n+1} + (1-\rho)\varepsilon_{\text{opt}}^{n}
\end{aligned}
$$

We can combine the above equations in two steps. First, we re-write the recurrence relations as:

$$
\varepsilon_{\text{pol}}^{n+2} = \varepsilon_{\text{pol}}^{n+1}(1-\pi) + \varepsilon_{\text{opt}}^{n+1}\pi\gamma \tag{12a}
$$

$$
\varepsilon_{\text{opt}}^{n+1} = \varepsilon_{\text{pol}}^{n} - \varepsilon_{\text{pol}}^{n+1} + (1-\rho)\varepsilon_{\text{opt}}^{n} \tag{12b}
$$

$$
\varepsilon_{\text{opt}}^{n}\pi\gamma = \varepsilon_{\text{pol}}^{n+1} - \varepsilon_{\text{pol}}^{n}(1-\pi) \tag{12c}
$$

By replacing $\varepsilon_{\text{pol}}^{n+1}$ in Eq. (12a) by Eq. (12b), and using Eq. (12c) to remove the $\varepsilon_{\text{opt}}^{n}$ term in the resulting equation, we obtain a second-order linear recurrence with constant coefficients for the policy error:

$$
\varepsilon_{\text{pol}}^{n+2} = ((1-\pi) + (1-\rho) - \pi\gamma)\varepsilon_{\text{pol}}^{n+1} + (\gamma\pi - (1-\pi)(1-\rho))\varepsilon_{\text{pol}}^{n}
$$

The determinant of the corresponding characteristic polynomial is

$$
\begin{aligned}
\Delta &= ((1-\pi) + (1-\rho) - \pi\gamma)^2 + 4(\gamma\pi - (1-\pi)(1-\rho)) \\
&= (\pi-\rho)^2 + \pi\gamma(\pi\gamma + 2\pi + 2\rho) \geq 0.
\end{aligned}
$$

Therefore, the roots of the polynomial are :

$$\begin{aligned}\beta_{\pm} &= \frac{(1-\pi)+(1-\rho)-\pi\gamma \pm \sqrt{\Delta}}{2} \\ &= 1 - \frac{\rho+\pi(1+\gamma)}{2}\left(1 - \pm\sqrt{\frac{\Delta}{(\rho+\pi(1+\gamma))^2}}\right) \\ &= 1 - \frac{\rho+\pi(1+\gamma)}{2}\left(1 - \pm\sqrt{\frac{(\pi(1+\gamma)+\rho)^2-4\pi\rho}{(\rho+\pi(1+\gamma))^2}}\right) \\ &= 1 - \frac{\rho+\pi(1+\gamma)}{2}\left(1 - \pm\sqrt{1-\frac{4\pi\rho}{(\rho+\pi(1+\gamma))^2}}\right).\end{aligned}$$

We therefore know that $\varepsilon_{\text{pol}}^n = A\beta_+^n + B\beta_-^n$. Now, we can use Eq. (12c) to obtain:

$$\begin{aligned}\varepsilon_{\text{opt}}^n &= \frac{1}{\pi\gamma}\left[A\beta_+^{n+1} + B\beta_-^{n+1} - A(1-\pi)\beta_+^n - B(1-\pi)\beta_-^n\right] \\ &= \frac{A(\beta_+ + \pi - 1)}{\pi\gamma}\beta_+^n + \frac{B(\beta_- + \pi - 1)}{\pi\gamma}\beta_-^n.\end{aligned}$$

Therefore, there exist constants $C$ and $D$ such that

$$\varepsilon_{\text{tot}}^n = \varepsilon_{\text{opt}}^n + \varepsilon_{\text{pol}}^n = C\beta_+^n + D\beta_-^n.$$

We need to find out if $C = 0$ (to know which root is the highest). We use the initial conditions:

$$\begin{aligned}\varepsilon_{\text{tot}}^0 &= \varepsilon_{\text{opt}}^0 + \varepsilon_{\text{pol}}^0 = 2 \\ \varepsilon_{\text{tot}}^1 &= \varepsilon_{\text{pol}}^0 + (1-\rho)\varepsilon_{\text{opt}}^0 = 2-\rho,\end{aligned}$$

and we deduce

$$\begin{aligned}&\varepsilon_{\text{tot}}^0 = 2 \implies C + D = 2 \\ &\varepsilon_{\text{tot}}^1 = (1-\pi)+(1-\rho)-\pi\gamma+\sqrt{\Delta}(C-1) \\ &\implies 2-\rho = (1-\pi)+(1-\rho)-\pi\gamma+\sqrt{\Delta}(C-1) \\ &\implies \pi(1+\gamma) = \sqrt{\Delta}(C-1).\end{aligned}$$

The left-hand side is non-negative, which means the right-hand side must verify $C - 1 \geq 0 \implies C \geq 1 > 0$. Therefore, we have $\varepsilon_{\text{tot}}^n = \Theta((1-\alpha)^n)$, with

$$\alpha = \frac{\rho+\pi(1+\gamma)}{2}\left(1 - \sqrt{1-\frac{4\pi\rho}{(\rho+\pi(1+\gamma))^2}}\right).$$

□

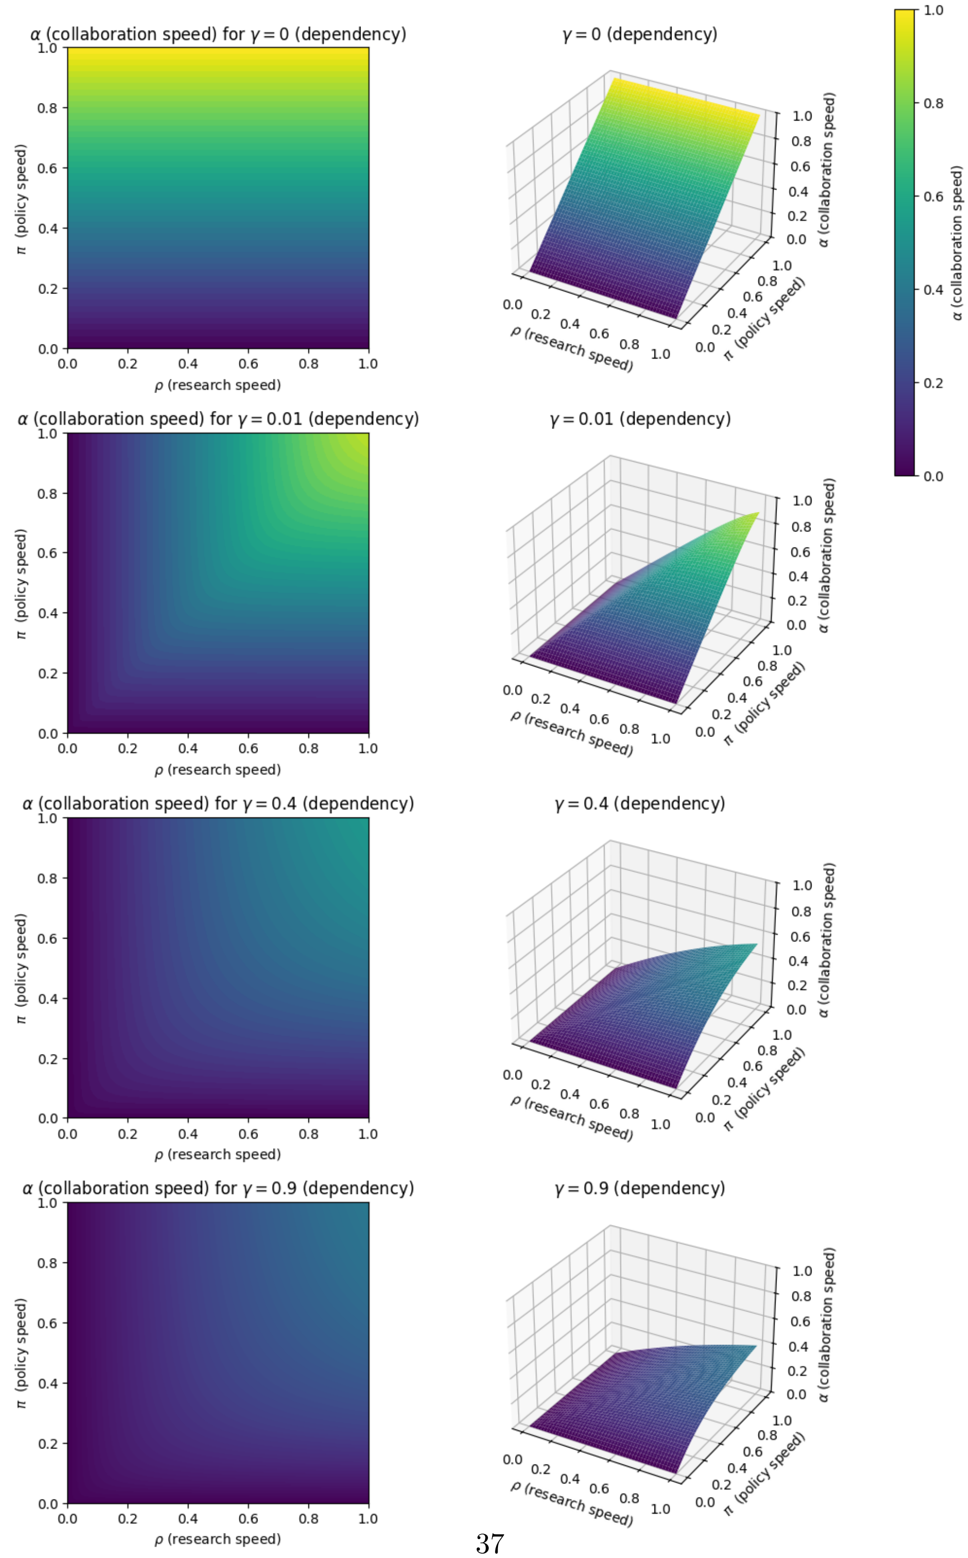




Figure 11: Convergence of policy-research collaboration.

*Note.* Illustrates the convergence rate $\alpha$ (lower is better) from Lemma 1, all figures present $\alpha$ as a function of $\pi$ and $\rho$ for a given value of $\gamma$.

*Proof.* Proof of Theorem 1. Recall that the convergence rate is given by $1-\alpha$, where

$$\alpha = \frac{\rho+\pi(1+\gamma)}{2}\left(1-\sqrt{1-\frac{4\pi\rho}{(\rho+\pi(1+\gamma))^2}}\right).$$

**Effect of $\gamma$.** We would like to show that the partial derivative with respect to $\gamma$ is non-positive. We first define $E=\rho+\pi(1+\gamma)$ and observe that $\partial E/\partial\gamma=\pi$. We can re-write $\alpha$ as

$$\alpha = \frac{E}{2}\left(1-\sqrt{1-\frac{4\pi\rho}{E^2}}\right) = \frac{E}{2}-\frac{1}{2}\sqrt{E^2-4\pi\rho}.$$

We can then compute the derivative as:

$$\begin{aligned}\frac{\partial\alpha}{\partial\gamma} &= \frac{1}{2}\frac{\partial E}{\partial\gamma}-\frac{1}{2}\frac{1}{2}\frac{2\frac{\partial E}{\partial\gamma}E}{\sqrt{E^2-4\pi\rho}}\\ &= \frac{\pi}{2}-\frac{\pi}{2}\frac{E}{\sqrt{E^2-4\pi\rho}}\\ &= \frac{\pi}{2}\left(1-\sqrt{\frac{E^2}{E^2-4\pi\rho}}\right)\end{aligned}$$

The partial derivative is well-defined and non-positive for $\pi,\rho\geq 0$, so $\alpha$ is decreasing in $\gamma$.

**Effect of $\pi$.** We can analogously compute the derivative as:

$$\begin{aligned}\frac{\partial\alpha}{\partial\pi} &= \frac{1}{2}\frac{\partial E}{\partial\pi}-\frac{1}{2}\frac{1}{2}\frac{2\frac{\partial E}{\partial\pi}E-4\rho}{\sqrt{E^2-4\pi\rho}}\\ &= \frac{1+\gamma}{2}-\frac{1}{2}\frac{(1+\gamma)E-2\rho}{\sqrt{E^2-4\pi\rho}}\\ &= \frac{1+\gamma}{2}\left(1-\frac{E-\frac{2\rho}{1+\gamma}}{\sqrt{E^2-4\pi\rho}}\right)\end{aligned}$$

The partial derivative is non-negative as long as

$$1-\frac{E-\frac{2\rho}{1+\gamma}}{\sqrt{E^2-4\pi\rho}}\geq 0 \Leftrightarrow E-\frac{2\rho}{1+\gamma}\leq\sqrt{E^2-4\pi\rho}.$$

If $\pi\leq\rho\frac{1-\gamma}{(1+\gamma)^2}$, the left-hand side is negative and the inequality trivially holds. Otherwise, the inequality is equivalent to

$$E^2+\frac{4\rho^2}{(1+\gamma)^2}-\frac{4\rho}{1+\gamma}E\leq E^2-4\pi\rho \Leftrightarrow \frac{\rho}{(1+\gamma)^2}-\frac{\pi(1+\gamma)+\rho}{1+\gamma}+\pi\leq 0 \Leftrightarrow \frac{\rho\gamma}{(1+\gamma)^2}\geq 0,$$

which is true when $\gamma, \rho \geq 0$. Therefore, $\alpha$ is non-decreasing in $\pi$.

**Effect of $\rho$.** We can analogously compute the derivative as:

$$\begin{aligned}
\frac{\partial \alpha}{\partial \rho} &= \frac{1}{2}\frac{\partial E}{\partial \rho} - \frac{1}{2}\frac{1}{2}\frac{2\frac{\partial E}{\partial \rho}E - 4\pi}{\sqrt{E^2 - 4\pi\rho}} \\
&= \frac{1}{2} - \frac{1}{2}\frac{E - 2\pi}{\sqrt{E^2 - 4\pi\rho}} \\
&= \frac{1}{2}\left(1 - \frac{E - 2\pi}{\sqrt{E^2 - 4\pi\rho}}\right)
\end{aligned}$$

The partial derivative is non-negative as long as

$$1 - \frac{E - 2\pi}{\sqrt{E^2 - 4\pi\rho}} \geq 0 \Leftrightarrow E - 2\pi \leq \sqrt{E^2 - 4\pi\rho}.$$

If $\pi \geq \rho + \gamma\pi$, the left-hand side is negative and the inequality trivially holds. Otherwise, the inequality is equivalent to

$$E^2 + 4\pi^2 - 4\pi E \leq E^2 - 4\pi\rho \Leftrightarrow E - \rho - \pi \geq 0 \Leftrightarrow \gamma \geq 0.$$

Therefore the derivative is non-negative and $\alpha$ is non-decreasing in $\rho$.

**Limit.** Furthermore, in the limit as $\gamma$ tends to zero, we can write:

$$\begin{aligned}
\lim_{\gamma \to 0} \alpha &= \frac{\rho + \pi}{2}\left(1 - \sqrt{1 - \frac{4\pi\rho}{(\rho + \pi)^2}}\right) \\
&= \frac{\pi + \rho}{2}\left(1 - \sqrt{\left(\frac{\pi - \rho}{\pi + \rho}\right)^2}\right) \\
&= \frac{\pi + \rho}{2} - \frac{1}{2}\sqrt{(\pi - \rho)^2} \\
&= \frac{\max(\pi, \rho) + \min(\pi, \rho)}{2} - \frac{1}{2}\left(\max(\pi, \rho) - \min(\pi, \rho)\right) \\
&= \min(\pi, \rho).
\end{aligned}$$

□

*Proof.* Proof of Theorem 2. As in the proof of Theorem 1, we once again define $E = \rho + \pi(1 + \gamma)$ so that we can write the convergence rate as

$$\alpha = \frac{E}{2}\left(1 - \sqrt{1 - \frac{4\pi\rho}{E^2}}\right) = \frac{E}{2} - \frac{1}{2}\sqrt{E^2 - 4\pi\rho}.$$

Let $k = \frac{\rho}{\pi}$, such that $\rho = k\pi$. We can then re-write $E = \pi(k+\gamma+1)$, such that $\partial E/\partial k = \pi$, and the convergence rate as

$$\alpha = \frac{E}{2}\left(1 - \sqrt{1 - \frac{4\pi\rho}{E^2}}\right) = \frac{E}{2} - \frac{1}{2}\sqrt{E^2 - 4k\pi^2}.$$

From the proof of Theorem 1, we know that

$$\frac{\partial \alpha}{\partial \gamma} = \frac{\pi}{2}\left(1 - \sqrt{\frac{E^2}{E^2 - 4k\pi^2}}\right),$$

and we can analogously take the derivative with respect to $k$,

$$\begin{aligned}
\frac{\partial \alpha}{\partial k} &= \frac{1}{2}\frac{\partial E}{\partial k} - \frac{1}{2}\frac{\partial}{\partial k}\left[\left(E^2 - 4k\pi^2\right)^{\frac{1}{2}}\right] \\
&= \frac{\pi}{2} - \frac{1}{2}\frac{1}{2}\frac{\partial}{\partial k}\left(E^2 - 4k\pi^2\right)\frac{1}{\sqrt{E^2 - 4k\pi^2}} \\
&= \frac{\pi}{2} - \frac{1}{4}(2\pi E - 4\pi^2)\frac{1}{\sqrt{E^2 - 4k\pi^2}} \\
&= \frac{\pi}{2}\left(1 - \frac{E - 2\pi}{\sqrt{E^2 - 4k\pi^2}}\right)
\end{aligned}$$

Then we can write

$$\begin{aligned}
\frac{\partial \alpha}{\partial \gamma} + \frac{\partial \alpha}{\partial k} \le 0 &\Leftrightarrow \frac{\pi}{2}\left(\sqrt{\frac{E^2}{E^2 - 4k\pi^2}} - 1\right) + \frac{\pi}{2}\left(\frac{E - 2\pi}{\sqrt{E^2 - 4k\pi^2}} - 1\right) \ge 0 \\
&\Leftrightarrow \frac{E}{\sqrt{E^2 - 4k\pi^2}} + \frac{E - 2\pi}{\sqrt{E^2 - 4k\pi^2}} \ge 2 \\
&\Leftrightarrow \frac{E - \pi}{\sqrt{E^2 - 4k\pi^2}} \ge 1 \\
&\Leftrightarrow E^2 - 2\pi E + \pi^2 \ge E^2 - 4k\pi^2 \\
&\Leftrightarrow \pi^2(1 + 4k) \ge 2\pi^2(k + \gamma + 1) \\
&\Leftrightarrow k \ge \frac{\gamma}{3}.
\end{aligned}$$

□

# B Formulation extensions and transportation simulation details

## B.1 Incorporating early release schedules in the optimization model

It is common for individual schools to have early release in the afternoon once a week (e.g. on Tuesday) for extracurricular activities. We can model this case as having two distinct afternoon routing scenarios: one for Tuesday, and one for all other weekdays. Early release leads to a different school day length and impacts the distribution of active PM runs. Depending on the context, the decision maker may care about one scenario more than another. Thus, we define $q_k$ as the weight that scenario $k$ takes in the final optimization problem.

We now give an example of the tensor formulation in the case of early release. Suppose there are five scenarios for school $s$, each representing a school day, i.e. $\mathcal{K} = \{\text{Mon}, \text{Tue}, \text{Wed}, \text{Thu}, \text{Fri}\}$. Let $e_{s,k}$ represents the number of early release hours in scenario $k$; that is, if school $s$ releases one hour earlier on Wednesday than the normal school length $\ell_s$, then $e_{s,\text{Wed}} = 60$.

For AM runs, early release does not impact the transportation costs: the number of active runs in Equation (10a) does not depend on the length of the school day, and thus stays the same regardless of the scenario; in other words, $N^{\text{AM}}_{s,t,\theta,k} = N^{\text{AM}}_{s,t,\theta,\text{Mon}}$ for all $k \in \mathcal{K}$. However, the PM runs critically depend on the length of the school day; at time $\theta$, a PM run $r$ in scenario $k$ is active if

$$h_t + \ell_s - e_{s,k} \leq \theta \leq h_t + \ell_s - e_{s,k} + \delta_{s,r}, \tag{13}$$

obtained by replacing the length of day $\ell_s$ in Eq. (10b) with $(L_s - E_{s,r})$. As a result, the total number of active PM runs for school $s$ in scenario $r$ is given by

$$N_{s,t,\theta,k} = \sum_{r=1}^{R_s^{\text{PM}}} \mathbb{1}\left(h_t + \ell_s - e_{s,k} \leq \theta \leq h_t + \ell_s - e_{s,k} + \delta_{s,r}\right). \tag{14}$$

The transportation cost minimization problem can then be formulated as:

$$\min_{\boldsymbol{x}} \sum_{k \in \mathcal{K}} q_k B_k \tag{15a}$$

$$\text{s.t. } B_k \geq \sum_{s,t} x_{s,t}(N^{\text{AM}}_{s,t,\theta,k} + N^{\text{PM}}_{s,t,\theta,k}), \qquad \forall \theta \in \Theta, \forall k \in \mathcal{K} \tag{15b}$$

$$\sum_{t=1}^{T} x_{s,t} = 1, \qquad \forall s \in [S] \tag{15c}$$

Recall that $q_k$ represents the weight for scenario $k$; $B_k$ represents the total transportation cost in scenario $k$ under assignment $\boldsymbol{x}$, which is the maximum number of active runs (buses) throughout a day in scenario $k$ (see Eq. (15b)). Note that using a weighted average cost in Eq. (15a) as the objective implies that the school district can use a different number of buses on each day of the week without incurring additional costs; this assumes that the school district can negotiate with the bus provider to have additional buses only on certain days to accommodate early release programs. This was the case for SFUSD.

However, depending on the contract, it is also possible that some school districts are charged based on the maximum fleet size on any given day, even if some buses are not utilized every day. In that case, the maximum number of buses across all scenarios, $\max_{k \in \mathcal{K}} B_k$, is a better approximation of the transportation cost and remains linearizable using one additional auxiliary variable.

## B.2 Details of transportation simulations and additional results

**Transportation simulator.** We first describe our school transportation simulator in detail. The input of the simulator is a district, i.e., a set of school locations, and the locations of students attending that school, as well as a proposed start time for each school. The distance between any two points is given by the Haversine metric, and travel times are computed assuming a constant vehicle speed of 10 mph.

We then compute a feasible routing solution by using the standard decomposition of the school bus routing algorithm first introduced by Desrosiers et al. (1986) and described in detail by Park and B. I. Kim (2010): (i) bus stop assignment, in which students from the same school are clustered into stops; (ii) bus routing, in which stops for the same school are connected into runs; (iii) bus scheduling, in which consecutive runs for different schools are assigned to a bus. Our approach closely follows the school bus routing algorithm described in Bertsimas, Delarue, and Martin (2019), with a few key modifications.

For the stop assignment step, we formulate the following integer program for each school:

$$\begin{aligned}
\min \quad & \sum_{p \in \mathcal{P}} y_p && && \text{(16a)} \\
\text{s.t.} \quad & \sum_{p' \in \mathcal{P}} z_{p,p'} = 1 && \forall p \in \mathcal{P} && \text{(16b)} \\
& \sum_{p' \in \mathcal{P}} z_{p',p} \leq n_p y_p && \forall p \in \mathcal{P} && \text{(16c)} \\
& z_{p,p'} \leq y_{p'} && \forall p, p' \in \mathcal{P} && \text{(16d)} \\
& y_p \in \{0,1\} && \forall p \in \mathcal{P} && \text{(16e)} \\
& z_{p,p'} \in \{0,1\} && \forall p, p' \in \mathcal{P}. && \text{(16f)}
\end{aligned}$$

The set $\mathcal{P}$ designates the set of students (pupils) for the school under study. The binary variable $y_p$ takes the value 1 if the address of student $p$ is used as a bus stop, and 0 otherwise. The binary variable $z_{p,p'}$ takes the value 1 if student $p$ is assigned to the bus stop at the address of student $p'$. Constraint (16b) ensures every student is assigned to a stop. Constraint (16c) ensures that no more than $n_p$ students are assigned to the address of student $p$. Constraint (16d) ensures that student $p$ is not assigned to the address of student $p'$ unless that address is used as a bus stop. The objective is to minimize the number of bus stops. We take $n_p = 30$ for all schools and all addresses. We also take into account walking restrictions (specified in the dataset) by fixing some $z_{p,p'}$ variables to zero if they correspond to an assignment that exceeds the walking distance.

Given the set of bus stops for each school, we next connect stops into runs. This problem is a special kind of open capacitated vehicle routing problem, which we solve using the OR-Tools routing library (Furnon and Perron 2024). We impose a pickup time of 60s at each stop and 240s at each school, a maximum route length of either 65 or 70 minutes, and a capacity of 100 students. We instruct the solver to balance total route length and number of vehicles by setting a penalty cost per vehicle of either 0s or 100000s. The travel time between stops is given by the Haversine distance, assuming the vehicle travels at a constant speed of 10mph. We use the guided local search algorithm with a time limit of 5 minutes. For each school, we compute four sets of runs, corresponding to the possible maximum route lengths $\{65 \text{ min}, 70 \text{ min}\}$ and the bus penalty costs $\{0, 100000\}$. This follows the insight from Bertsimas, Delarue, and Martin (2019) that computing diverse sets of runs for each school, and selecting which set we use at scheduling time, can significantly reduce total fleet size.

We use the four sets of runs for each school when computing the run-counting tensor $\boldsymbol{N}$ for our formulation, as well as an input to the quadratic assignment formulation of Bertsimas, Delarue, and Martin (2019). In both cases, we average the number of runs obtained from each of the four scenarios.

Given runs for each school, and a set of start times produced by any algorithm, the final step in our simulation is to solve the bus scheduling problem. For this step, we use our implementation of the BiRD algorithm of Bertsimas, Delarue, and Martin (2019), which jointly selects the optimal set of runs for each school and connects runs into full bus itineraries spanning the entire morning. We provide our code for all start time optimization and school bus routing algorithms as a companion to this paper.

**Additional results.** We now detail additional results of interest from our simulation-based comparison of start time optimization algorithms. Bertsimas, Delarue, and Martin (2019) stated that their quadratic assignment algorithm includes some hyperparameters that need to be tuned for each district. We sought to quantify the need for parameter tuning, since it represents a significant amount of additional computations.

Figure 12 shows the impact of parameter tuning on the quadratic assignment algorithm of Bertsimas, Delarue, and Martin (2019). For small districts, the gain from tuning is minimal. However, for larger districts, and a larger gap between start time options, the

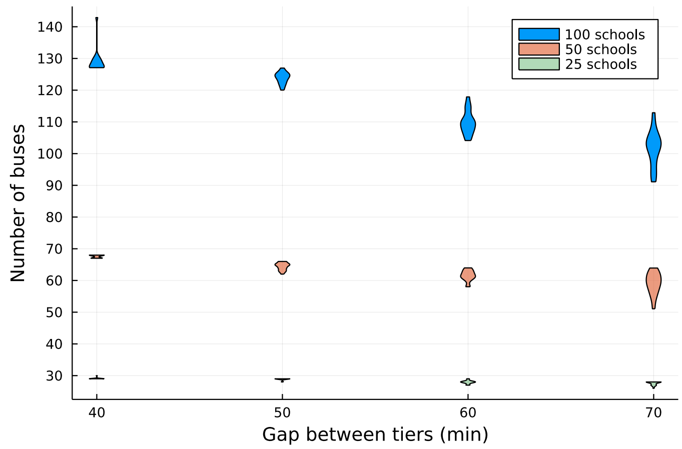


Figure 12: Effect of parameter tuning on the quadratic bell time optimization algorithm of Bertsimas, Delarue, and Martin 2019.
*Note.* In each of the 12 scenarios, we consider 16 hyperparameter configurations for the bell time optimization algorithm, with the downstream transportation cost of all configurations depicted as a violin plot. When the number of schools and the gap between tiers grows large, tuning can improve transportation costs by up to 20%.

difference between the worst parameter configuration and the best one can easily reach 20% — justifying the need for parameter tuning.

Finally, we take a deeper look at the solutions produced by our algorithm and the quadratic assignment formulation, in order to understand how they differ beyond the number of buses required. Figure 13 shows some justification for the performance difference between our formulation and the quadratic formulation from Bertsimas, Delarue, and Martin (2019). The quadratic approach puts much more emphasis on reducing the number of schools starting in the middle tier (to allow for more guaranteed connections). This works reasonably well when the gap between tiers is small: in this case, the linear formulation also seeks to have fewer schools in the middle tier, though this effect is less pronounced. However, when the gap between tiers is large, the linear and quadratic formulations give vastly different answers (even though, according to Figure 4, they have similar transportation costs). The quadratic formulation still prefers having fewer schools in the middle tier, while the linear formulation actually places the fewest schools in the earliest tier. This result also empirically confirms that there are many vastly different solutions with similar transportation costs.

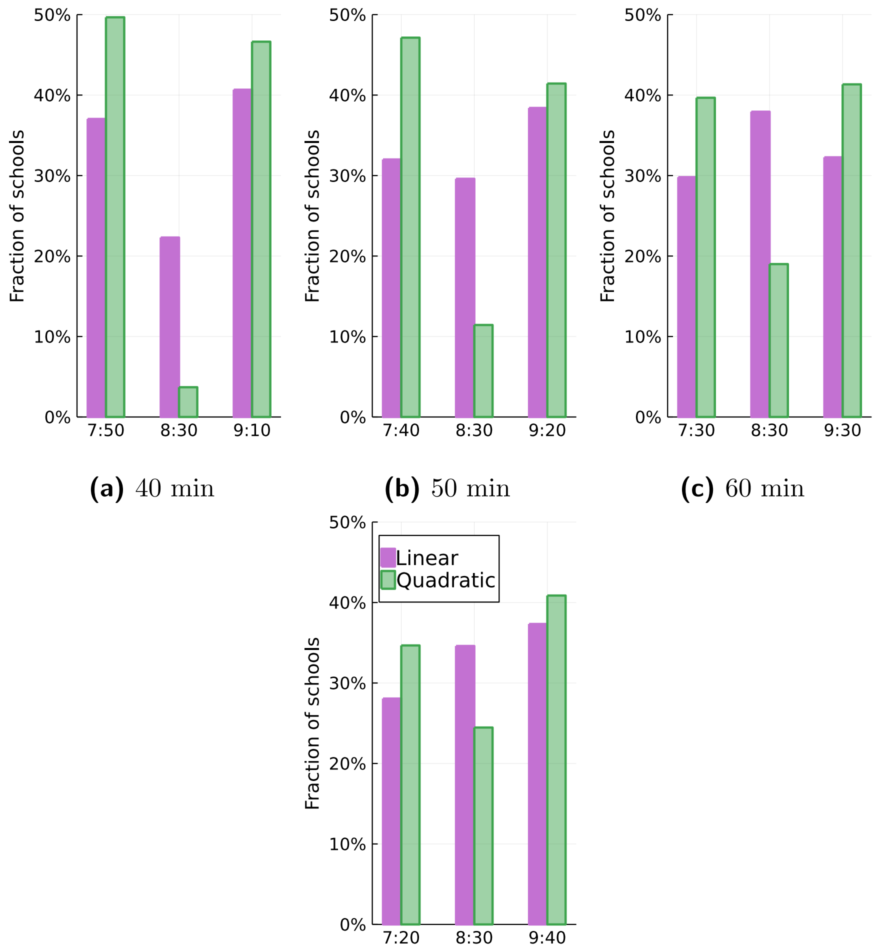


**(a)** 40 min **(b)** 50 min **(c)** 60 min

**(d)** 70 min

Figure 13: Distribution of start times for different optimization algorithms.
*Note.* We compare our linear tensor-based formulation to Bertsimas, Delarue, and Martin 2019's quadratic assignment formulation. The linear formulation is also considerably more likely to have schools starting in the middle tier, which is often the most desirable.

## C Comparing different solutions with similar costs

Table 2: Comparison of two start time schedules with similar costs for known objectives

| School Name | Schedule 1 | Schedule 2 |
|---|---|---|
| Giannini (A.P.) MS | 9:30 AM | 9:30 AM |
| Lincoln (Abraham) HS | 8:40 AM | 8:40 AM |
| Alamo ES | 8:40 AM | 8:40 AM |
| Alvarado ES | 7:50 AM | 7:50 AM |
| Aptos MS | 9:30 AM | 9:30 AM |
| Argonne ES | 8:40 AM | 8:40 AM |
| Balboa HS | 8:40 AM | 8:40 AM |
| Carmichael (Bessie) K-5 | 8:40 AM | 8:40 AM |
| Harte (Bret) ES | 8:40 AM | 8:40 AM |
| Bryant ES | 7:50 AM | 7:50 AM |
| Lee, Edwin and Anita Newcomer School | 8:40 AM | 8:40 AM |
| Clarendon ES | 9:30 AM | 9:30 AM |
| Lilienthal (Claire) K-2 Sacramento | 9:30 AM | 9:30 AM |
| Cleveland ES | 9:30 AM | 9:30 AM |
| Yu (Alice Fong) K-8 | 9:30 AM | 9:30 AM |
| Sloat (Commodore) ES | 8:40 AM | 8:40 AM |
| Lau (Gordon J) ES | 8:40 AM | 8:40 AM |
| SF Community K-8 (K-8) | 9:30 AM | 9:30 AM |
| Webster (Daniel) ES | **8:40 AM** | **7:50 AM** |
| Milk (Harvey) ES | 9:30 AM | 9:30 AM |
| Drew (Dr Charles) ES | 7:50 AM | 7:50 AM |
| CIS at DeAvila ES | 8:40 AM | 8:40 AM |
| Taylor (Edward R) ES | **8:40 AM** | **7:50 AM** |
| El Dorado ES | 7:50 AM | 7:50 AM |
| Cobb (Dr William L) ES | **7:50 AM** | **8:40 AM** |
| Everett MS | 9:30 AM | 9:30 AM |
| Dolores Huerta | 7:50 AM | 7:50 AM |
| Feinstein (Dianne) ES | 7:50 AM | 7:50 AM |
| Key (Francis Scott) ES | 7:50 AM | 7:50 AM |
| Francisco MS | 9:30 AM | 9:30 AM |
| McCoppin (Frank) PK | **8:40 AM** | **9:30 AM** |
| McCoppin (Frank) ES | **8:40 AM** | **9:30 AM** |
| Galileo HS | 8:40 AM | 8:40 AM |
| Garfield ES | 7:50 AM | 7:50 AM |
| Peabody (George) ES | 7:50 AM | 7:50 AM |
| Washington (George) HS | 8:40 AM | 8:40 AM |

| School Name | Schedule 1 | Schedule 2 |
|---|---|---|
| Glen Park ES | 7:50 AM | 7:50 AM |
| Grattan ES | 7:50 AM | 7:50 AM |
| Guadalupe ES | 8:40 AM | 8:40 AM |
| Chavez (Cesar) ES | **9:30 AM** | **8:40 AM** |
| Hoover (Herbert) MS | 9:30 AM | 9:30 AM |
| Hillcrest ES | 7:50 AM | 7:50 AM |
| Buena Vista Horace Mann 6-8 | 9:30 AM | 9:30 AM |
| SF International HS | 8:40 AM | 8:40 AM |
| Carver (Dr George W) ES | 8:40 AM | 8:40 AM |
| Denman (James) MS | 9:30 AM | 9:30 AM |
| Lick (James) MS | 9:30 AM | 9:30 AM |
| Parker (Jean) ES | 8:40 AM | 8:40 AM |
| Jefferson ES | 7:50 AM | 7:50 AM |
| Muir (John) ES | 8:40 AM | 8:40 AM |
| O'Connell (John) HS | 8:40 AM | 8:40 AM |
| Serra (Junipero) ES | 8:40 AM | 8:40 AM |
| Lafayette ES | 7:50 AM | 7:50 AM |
| Lafayette ES | 9:30 AM | 9:30 AM |
| Lakeshore ES | **7:50 AM** | **9:30 AM** |
| Lawton K-8 (K-5) | 9:30 AM | 9:30 AM |
| Flynn (Leonard R) ES | 8:40 AM | 8:40 AM |
| Longfellow ES | 8:40 AM | 8:40 AM |
| Lowell HS | 8:40 AM | 8:40 AM |
| Marina MS | 9:30 AM | 9:30 AM |
| King Jr (Dr Martin L) MS | 9:30 AM | 9:30 AM |
| Marshall ES | 8:40 AM | 8:40 AM |
| McKinley ES | 7:50 AM | 7:50 AM |
| Miraloma ES | 7:50 AM | 7:50 AM |
| Moscone (George R) ES | 7:50 AM | 7:50 AM |
| Mission Education Center ES | 9:30 AM | 9:30 AM |
| Mission HS | 8:40 AM | 8:40 AM |
| Monroe ES | 8:40 AM | 8:40 AM |
| New Traditions ES | 9:30 AM | 9:30 AM |
| Downtown High School | 9:30 AM | 9:30 AM |
| Ida B. Wells HS | 9:30 AM | 9:30 AM |
| Ortega (Jose) ES | 7:50 AM | 7:50 AM |
| Sunset ES | 8:40 AM | 8:40 AM |
| Jordan (June) HS | 8:40 AM | 8:40 AM |
| Revere (Paul) K-5 | 9:30 AM | 9:30 AM |
| Burton (Phillip & Sala) HS | 8:40 AM | 8:40 AM |

| School Name | Schedule 1 | Schedule 2 |
|---|---|---|
| Presidio MS | 9:30 AM | 9:30 AM |
| Stevenson (Robert Louis) ES | 7:50 AM | 7:50 AM |
| Wallenberg (Raoul) HS | 8:40 AM | 8:40 AM |
| Parks (Rosa) ES | 7:50 AM | 7:50 AM |
| Redding ES | 8:40 AM | 8:40 AM |
| Rooftop K-4 | 9:30 AM | 9:30 AM |
| Roosevelt MS | 9:30 AM | 9:30 AM |
| Wo (Yick) ES | 9:30 AM | 9:30 AM |
| SF Public Montessori ES | 8:40 AM | 8:40 AM |
| Asawa (Ruth) SOTA HS | 8:40 AM | 8:40 AM |
| Sanchez ES | 7:50 AM | 7:50 AM |
| Sheridan ES | 7:50 AM | 7:50 AM |
| Sherman ES | 7:50 AM | 7:50 AM |
| Malcolm X Academy ES | 8:40 AM | 8:40 AM |
| The Academy - SF @ McAteer HS | 8:40 AM | 8:40 AM |
| Spring Valley Science ES | 8:40 AM | 8:40 AM |
| King (Thomas Starr) ES | **9:30 AM** | **7:50 AM** |
| Sunnyside ES | 7:50 AM | 7:50 AM |
| Sutro ES | 8:40 AM | 8:40 AM |
| Marshall (Thurgood) HS | 8:40 AM | 8:40 AM |
| Brown Jr. (Willie) MS | 9:30 AM | 9:30 AM |
| Tenderloin ES | **7:50 AM** | **8:40 AM** |
| Ulloa ES | 9:30 AM | 9:30 AM |
| Visitacion Valley ES | 8:40 AM | 8:40 AM |
| Visitacion Valley MS | 9:30 AM | 9:30 AM |
| Chin (John Yehall) ES | 9:30 AM | 9:30 AM |
| West Portal ES | 8:40 AM | 8:40 AM |
| Argonne (PreK-TK & OST) | 8:40 AM | 8:40 AM |
| Carmichael (Bessie) PK | 8:40 AM | 8:40 AM |
| L. Havard PK | 9:30 AM | 9:30 AM |
| Bryant PK | 7:50 AM | 7:50 AM |
| Las Americas PK | 7:50 AM | 7:50 AM |
| Mahler, Theresa S. Early Education School | 9:30 AM | 9:30 AM |
| Noriega EES | 8:40 AM | 8:40 AM |
| Sheridan PK | 7:50 AM | 7:50 AM |
| Grattan PK | 7:50 AM | 7:50 AM |
| Chavez (Cesar) PK | **9:30 AM** | **8:40 AM** |
| Ortega (Jose) PK | 7:50 AM | 7:50 AM |
| Jefferson PK | 8:40 AM | 8:40 AM |
| McLaren TK | 9:30 AM | 9:30 AM |

| School Name | Schedule 1 | Schedule 2 |
|---|---|---|
| Malcolm X Academy PK | 8:40 AM | 8:40 AM |
| Muir (John) PK | 8:40 AM | 8:40 AM |
| Revere PK | 9:30 AM | 9:30 AM |
| Presidio Early Education School | 9:30 AM | 9:30 AM |
| R. Weill PK | 7:50 AM | 7:50 AM |
| San Miguel PK | 9:30 AM | 9:30 AM |
| Cobb (Dr William L) PK | **7:50 AM** | **8:40 AM** |
| Tule Elk Park Early Education School | **9:30 AM** | **8:40 AM** |
| Carmichael (Bessie) 6-8 | 9:30 AM | 9:30 AM |
| Rooftop 5-8 | 9:30 AM | 9:30 AM |
| Lilienthal (Claire) 3-8 Divisadero | 9:30 AM | 9:30 AM |
| Yu (Alice Fong) K-8 (K-5) | 9:30 AM | 9:30 AM |
| SF Community K-8 (6-8) | 9:30 AM | 9:30 AM |
| Lawton K-8 (6-8) | 9:30 AM | 9:30 AM |
| Lau (Gordon) PreK | 8:40 AM | 8:40 AM |
| Buena Vista Horace Mann K-5 | 9:30 AM | 9:30 AM |
| Revere (Paul) 6-8 | 9:30 AM | 9:30 AM |

# D Sensitivity analysis tool

The sensitivity analysis tool is presented diagrammatically in Figure 14. Its core functionality leverages the fact that the tensor formulation of the transportation cost can be computed efficiently in real time by a simple spreadsheet model. As a result, policymakers can manually edit the start time of each school and observe the cost implications. On one hand, this second tool does not adjust the start times at other schools in order to keep cost and change low; as a result, solution performance can quickly deteriorate if too many adjustments are made. On the other hand, it truly gives policymakers the ability to explore, understand and select any solution they want.

The sensitivity analysis tool also provides additional transparency to policymakers, allowing them to peer into the black box of transportation cost modeling. Step 2 in Figure 14 illustrates how the transportation cost implications of schedule changes for the morning and afternoon are displayed to policymakers. The visualization clearly reveals why certain days are good or bad candidates for aligned early release. Coordinating end times can in general have negative effects on transportation since it creates a short window of acute bus need. However, Figure 14 shows that this is not always the case. Indeed, section B of the diagram shows that the change in transportation costs when adding an early release at 1 PM for all elementary schools is negligible because the system is already saturated in the morning.

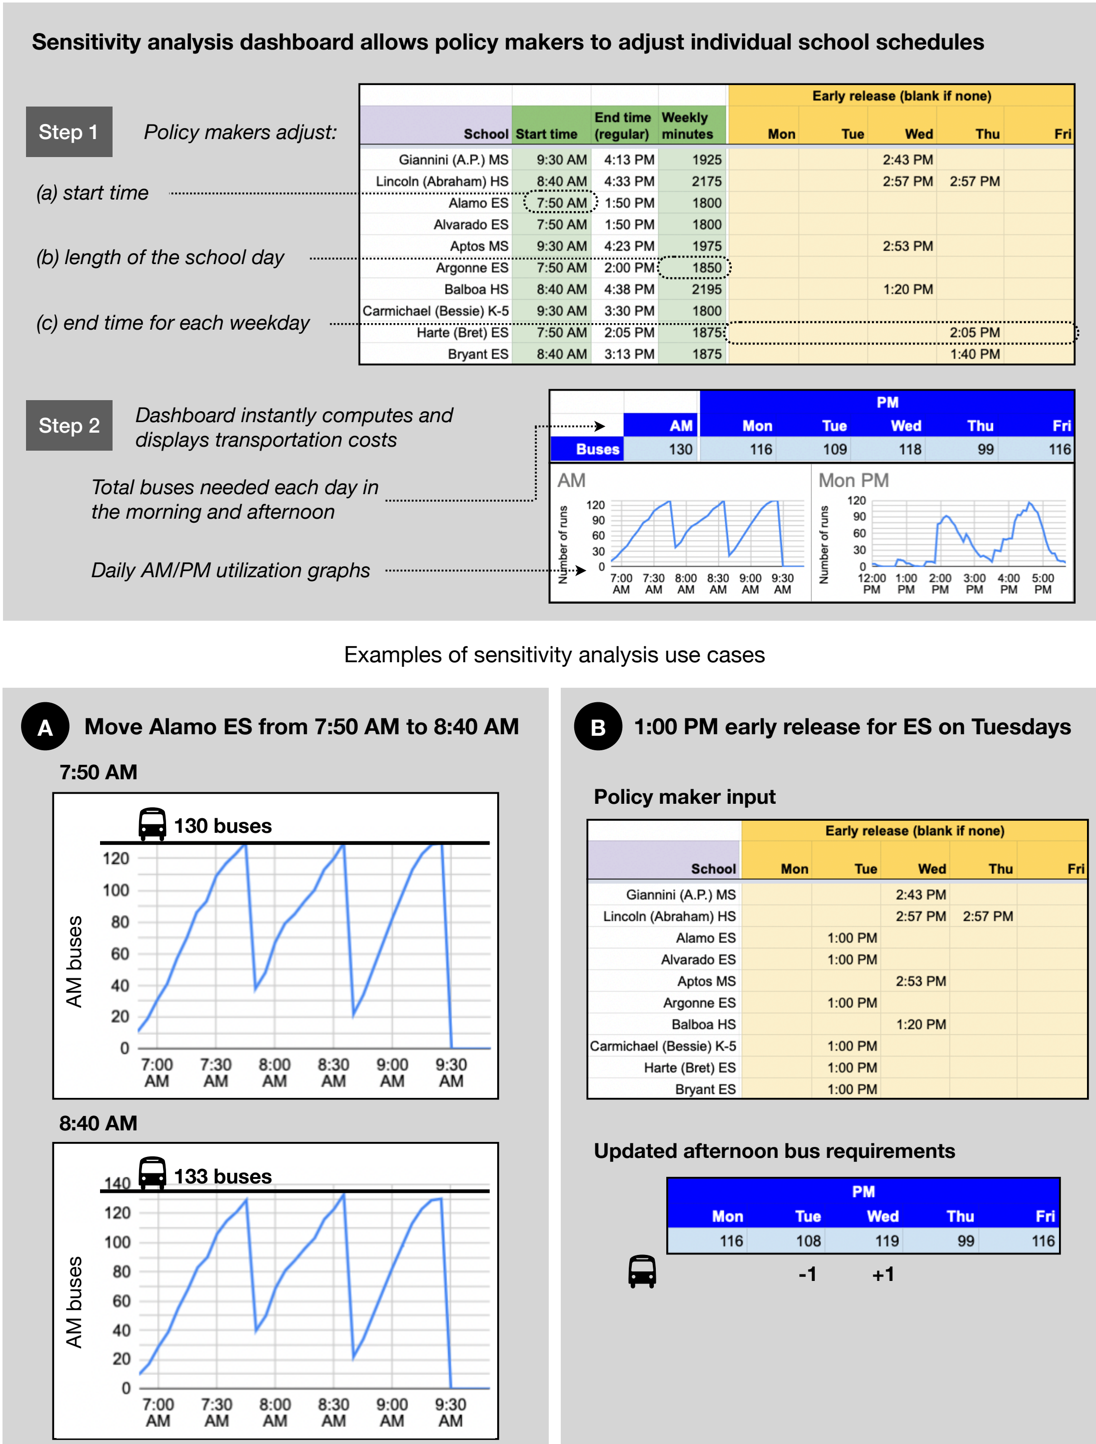


Figure 14: Diagram of the sensitivity analysis tool.

*Note.* Top section depicts a simplified view with editable fields. Box A and B illustrate two use cases.

## E Additional survey results

### E.1 Staff satisfaction of the schedule change

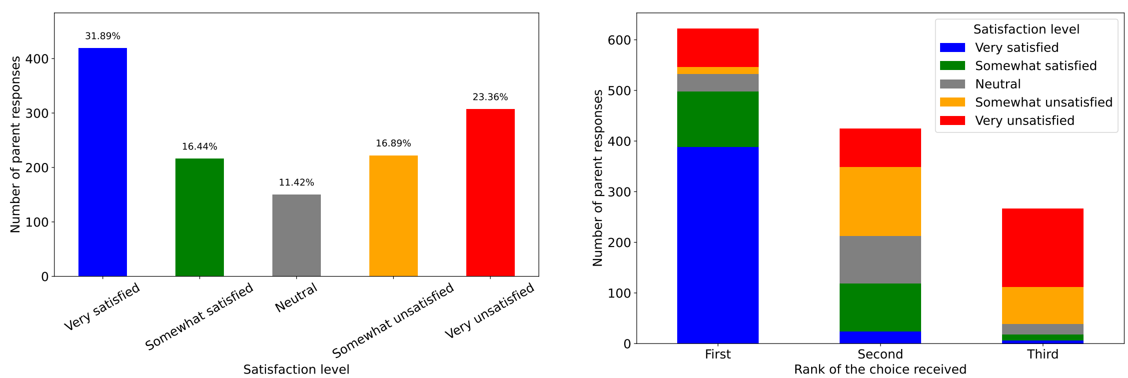


**(a)** Staff satisfaction by the satisfaction level **(b)** Staff satisfaction by rank of the choice received

Figure 15: Staff satisfaction level for the schedule, after the change

Figure 15 displays the satisfaction levels among SFUSD staff members, as well as their preference rankings over all three feasible start times, segmented by satisfaction level. This is similar to the information shown for parents in Figure 9 in Section 5.3. According to Figure 15a, approximately half of the responding staff members were at least somewhat satisfied with the final schedule. However, 23.36% reported being "very unsatisfied," highlighting the significant challenges and impacts faced by the staff. Conversely, Figure 15b reveals that a notable number of those who were "very unsatisfied" had actually received their top choice. This underscores the complexity of changing start times, indicating that a single solution may not be universally applicable.

### E.2 Survey screenshots

This section presents the survey questions that correlate with the statistics in Section 5.3. The survey was conducted using Qualtrics. It begins with a consent form, depicted in Figure 16. Subsequently, the survey includes five questions related to start times, followed by five demographic questions. These demographic questions cover the respondents' (or their children's) mode of transportation, race, ethnicity, and whether they receive English language development services, among other aspects. Of these questions, only the first two related to start times are discussed in this paper in Section 5.3 and illustrated in Figure 17. A complete set of survey questions is available upon request.

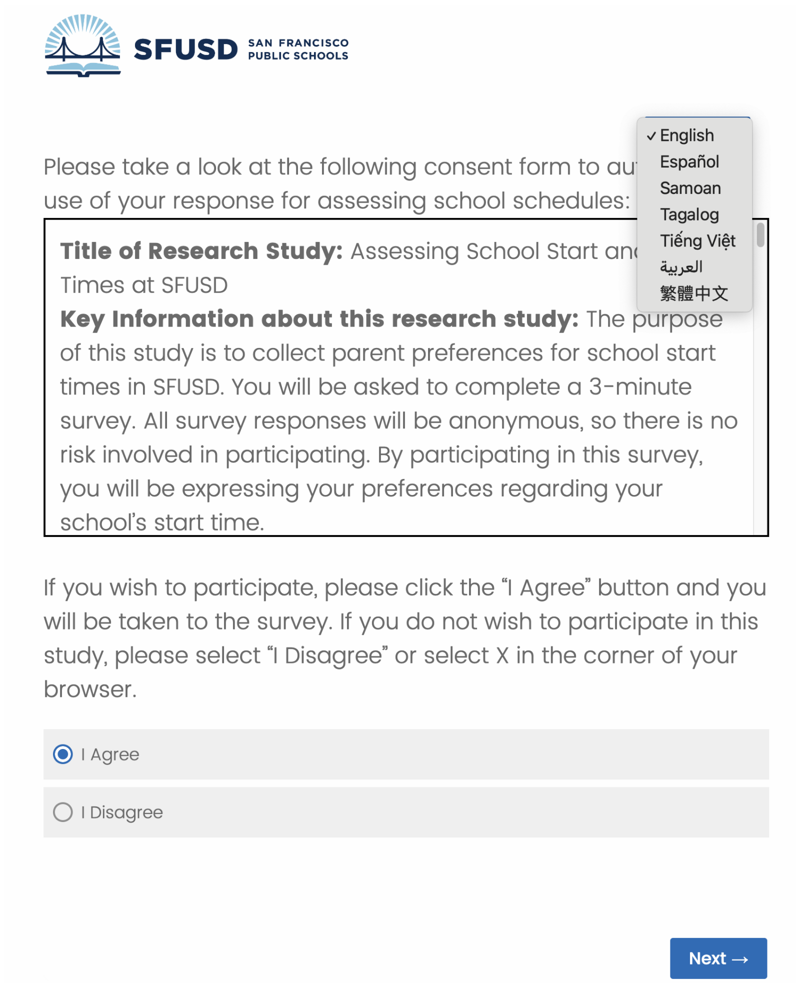


Figure 16: Survey consent form.

*Note.* The survey was offered in 7 languages (English, Spanish, Samoan, Tagalog, Vietnamese, Arabic, and Mandarin) to include parents who do not speak English as their primary language.

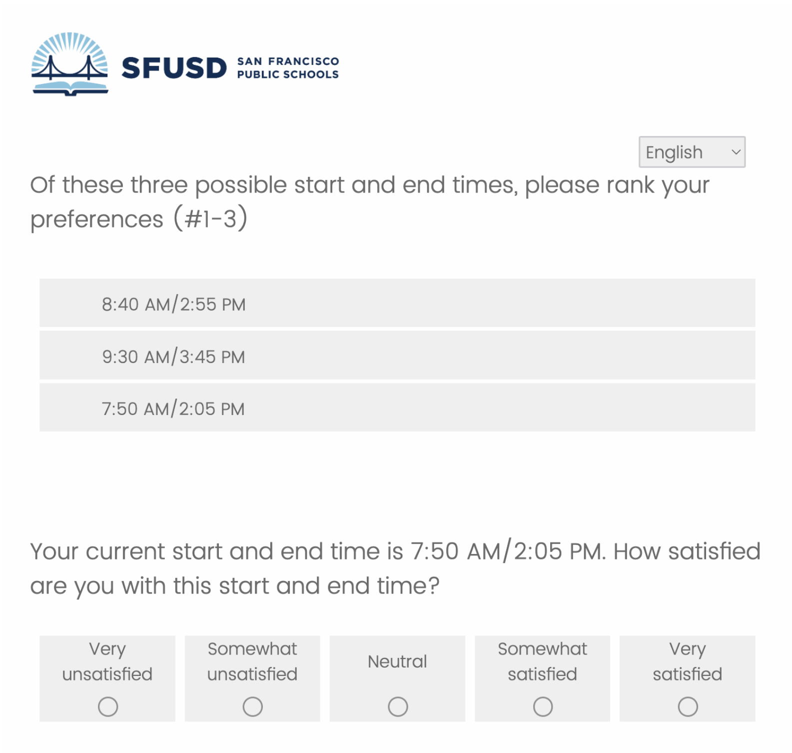

SFUSD SAN FRANCISCO PUBLIC SCHOOLS

English

Of these three possible start and end times, please rank your preferences (#1-3)

8:40 AM/2:55 PM

9:30 AM/3:45 PM

7:50 AM/2:05 PM

Your current start and end time is 7:50 AM/2:05 PM. How satisfied are you with this start and end time?

| Very unsatisfied | Somewhat unsatisfied | Neutral | Somewhat satisfied | Very satisfied |
|---|---|---|---|---|
| ○ | ○ | ○ | ○ | ○ |

Figure 17: Survey page for the rank of preferences and satisfaction.
*Note.* Respondents were asked to rank their preferences (via a drag-and-drop list) as well as indicate their satisfaction with the current start time. The survey form automatically fills in the correct current time based on the respondent's school.