Extending the Minimal Perturbation Problem with Disruptiveness Weighting for Dynamic University Timetabling

Helen Kirby-Hawkins

Manchester Metropolitan University, Manchester, UK
hkirbyhawkins@gmail.com

Abstract. University course timetabling is a dynamic optimization problem in which published schedules frequently require modification due to changes in enrolment, staffing, or resource availability. The Minimal Perturbation Problem (MPP) addresses this by seeking revised timetables that minimize deviation from an existing solution. While many standard formulations of MPP incorporate weighted measures of perturbation, these approaches do not typically account for how disruption varies with the timing of a change within the teaching period. This paper proposes a behaviorally informed extension to MPP that introduces a disruptiveness weighting for post-publication timetable changes. The proposed Minimally Disruptive Perturbation (MDP) model weights post-publication changes according to the week of occurrence, the type of change, and the number of students affected. This enables a more nuanced assessment of disruption in dynamic timetabling contexts. The model incorporates temporal weighting, change type and student reach to assess the behavioral cost, by way of a disruptiveness score, of timetable volatility. Using institutional data from a large UK university, the paper demonstrates how the disruptiveness score differentiates between changes that are operationally similar but behaviorally distinct. The approach is intended as a decision-support metric that can be incorporated into post-publication rescheduling processes and candidate solution evaluation, supporting more human-centric timetable optimization.

Keywords: *University timetabling, minimal perturbation problem, dynamic scheduling, disruptiveness weighting, human-centric optimization.*

1 Introduction

University course timetabling is commonly formulated as a constraint-based optimization problem in which limited resources—rooms, staff, and time slots—must be allocated while satisfying hard and soft constraints. While much of the literature focuses on the generation of feasible initial timetables, in practice university timetables remain subject to frequent change after publication. Fluctuations in student enrolment, late staffing changes, and external constraints often render published timetables infeasible, requiring re-optimization post-publication throughout the teaching semester.

The Minimal Perturbation Problem (MPP) was introduced to address this challenge by prioritizing solutions that minimize deviation from an existing timetable when changes are required. MPP has been widely adopted in dynamic timetabling environments and is typically operationalized by minimizing the number of changes or constraint violations relative to the original solution. While this approach is effective in preserving structural similarity, it assumes that all changes are equally disruptive, thereby overlooking their behavioral impact. In operational university settings, this assumption is problematic. A room change late in the semester is treated equivalently to a day or time change in the first week of teaching, despite the latter often having greater consequences for students and staff. Timetabling practitioners frequently recognize this distinction informally, but it is rarely reflected in optimization models or decision-support tools. More fundamentally, this approach implicitly privileges the existing timetable by penalizing deviation from it, effectively treating the published solution as a desirable reference state, even where it may not represent an optimal or stable schedule.

This paper proposes an extension to the MPP framework that incorporates disruptiveness weighting into post-publication timetable changes. The proposed Minimally Disruptive Perturbation (MDP) model introduces a quantitative metric that weighs changes according to their timing, type, and scale of impact. The contribution of this paper is not the introduction of weighting per se, but the explicit integration of temporal sensitivity into the evaluation of post-publication timetable changes, demonstrating how disruption varies with timing, type, and scale.

2 Related Work

2.1 University Timetabling and Scheduling

University course timetabling is a well-established problem in operations research, often framed as a constraint-based optimization challenge. The goal is to allocate limited resources — such as rooms, instructors, and time slots — while satisfying a set of hard and soft constraints (Lindahl et al., 2018).

Automated scheduling systems have become increasingly sophisticated, employing techniques such as integer programming (Phillips et al., 2017), genetic algorithms (Abdelhalim and El Khayat, 2016), and satisfiability solvers (Lemos et al., 2020) to generate feasible timetables.

Many real-life timetabling problems can be codified as constraint-based optimization problems. Tolerating dynamism in the university timetable is systematically managed by algorithms solving optimization problems, which are subject to disruptions that make the original solution no longer feasible, so the problem becomes dynamic and requires re-solving with minimal perturbations (Lemos et al., 2020).

However, these models typically prioritize computational efficiency and feasibility over human experience. The Minimal Perturbation Problem (MPP), introduced by Müller et al. (2005), addresses the challenge of modifying an existing timetable with minimal disruption. While MPP has been widely adopted in dynamic scheduling environments, it primarily focuses on minimizing the number (count) of changes or constraint violations, without considering their disruptive weight and effect on objects in the timetable (staff, students and resources). There is an anecdotal belief that not all types of changes or constraints are equally weighted. Many models account for hard and soft constraints, but these are not sufficiently nuanced to reflect weighting or the human behavioral impact of changes to the timetable when evaluating the best fit of a candidate option as a potential solution to the timetabling problem. Unlike soft constraint weighting within an optimization model, the proposed disruptiveness score operates as an evaluative metric applied to candidate post-publication solutions, rather than directly influencing solver search.

However, empirical and computational studies demonstrate that strict minimization of perturbations can lead to low quality solutions, and that allowing a modest increase in change often yields substantially improved outcomes (Phillips et al., 2017). Moreover, MPP formulations require institution-specific weighting policies that are external to the model itself (Kingston, 2016), and many approaches struggle to efficiently accommodate common real-world disruptions such as domain or capacity changes (Banbara et al., 2019). Collectively, this body of work suggests that classical MPP formulations prioritize structural similarity over the qualitative and behavioral impact of timetable changes on end users, motivating the need for richer, human-centered disruption models in university contexts. There also remains a gap in the literature regarding how timetabling changes may causally affect student behavior, particularly in the early weeks of a semester when new routines are still forming.

Recent work has highlighted the importance of human-centric consideration in timetabling, recognizing that schedules are consumed by real students — human end-users with complex lives — rather than abstract "student" objects, of which the literature speaks in somewhat reductionistic terms. Lindahl et al. (2018) advocate for a strategic view of timetabling that includes institutional goals and user preferences.

It is worth situating this contribution precisely against existing weighted-perturbation approaches, since weighting change severity is not itself new: El Sakkout et al. (1998) introduced distance-based measures of perturbation in dynamic scheduling, Müller et al. (2005) and the resulting UniTime system support a weighted sum of criteria including the number of affected students and the type of change, and several institution-specific MPP deployments already apply some form of change-type weighting in practice (Kingston, 2016). What these existing weighting schemes do not incorporate is an explicit temporal term: a given change type and student count are weighted identically regardless of whether the change occurs in week one or week eight of teaching. The contribution of this paper is therefore not the use of weighting per se, but the introduction of a temporal-sensitivity term (W) that decays with elapsed teaching time, combined multiplicatively with type and reach so that early, compound, high-reach changes are penalized disproportionately rather than additively. This paper further grounds the type and reach components empirically against institutional change data and student-reported disruption rankings from a single UK university, rather than relying solely on practitioner judgement, although the specific weight values themselves remain provisional (Sect. 4.3). Existing weighted approaches therefore capture structural differences between solutions, but do not explicitly model how the context in which a change is experienced alters its practical impact. These approaches therefore measure perturbation as a property of the solution structure, rather than as a function of when and how changes are experienced by end-users.

This study contributes to this discourse by extending the MPP framework and principles to consider behavioral dimensions, offering a disruption score that quantifies the impact of changes based on timing, type, and student reach. This formal weighting mechanism is reflective of the practical distinctions made between types of changes made by practitioners and offers a supportive model for decision-making.

3 Problem Context and Data

The study is situated within a large UK university operating a centrally managed timetable and was conducted between December 2024 and March 2025. Post-publication timetable changes occur frequently due to fluctuations in student numbers, module enrolment changes, staff amendments and room changes. Institutional timetable change data were extracted for a single academic semester, recording the week of change, the type of change, and the number of students affected. The analysis focuses on changes occurring within the first eight teaching weeks, where timetable volatility is highest. The dataset included all students on activities affected by two or more timetable changes — excluding minor administrative edits such as a change of named staff member or a cosmetic relabeling of an activity's display name, neither of which affects when, where, or how a student attends — in the first four weeks of teaching (N = 10,341). Student instances in this dataset do not correspond to unique students, as individual students may be affected by multiple timetable changes.

3.1 Quantitative Findings

Analysis of the institutional data reveals that 49.4% of the entire student population experienced at least one impactful timetable change within the first eight weeks of teaching. This equates to 21,991 unique students, with 11,852 experiencing two or more changes. A total of 906 changes and 747 additions or deletions were recorded during the period, affecting 46,940 student instances (acknowledging that unique students in this count will have been affected more than once, hence the total exceeding the population).

During the first four weeks of teaching, 10,341 students were impacted by timetable changes. These early disruptions are particularly significant given the importance of this period for habit formation.

Activity and Course-Level Disruption Analysis. To further analyze the nature of timetable volatility, an analysis was conducted on the types of activities and which courses were most frequently affected by timetable changes, drawn from the same institutional dataset. Using the activity naming conventions available in the dataset, activity types were examined from the "activity name" field. The most frequently changed activity types were:

Table 1. Total changes by activity type, number of changes in weeks 1–4, and percentage of total changes occurring in weeks 1–4, by activity type.

Activity Type	Total Changes	Changes in Weeks 1–4	% in Weeks 1–4
Seminar	205	153	74.63%
Lecture	192	124	64.58%
Tutorial	114	88	77.19%
PC session	111	45	40.54%
Workshop	91	70	76.92%
Practical	36	23	63.89%
Drop-In	28	22	78.57%
Web	20	9	45.00%
Groupwork	15	13	86.67%
Peer Assisted Learning (PAL)	13	13	100.00%

The table illustrates that the most frequently changed activity types are Seminars and Lectures. However, in weeks 1–4 the activity types with the highest percentage of early-week changes are PAL sessions and Groupwork. In Fig. 1 we can see the frequency of different types of change, including compound changes (such as day + time + room together), across the teaching weeks.

It should be noted that Table 1 reports raw change counts rather than change rates, and so the prominence of Seminars and Lectures partly reflects the fact that these are the most common activity types in the timetable overall, rather than indicating that these activity types are disproportionately prone to change. A more informative comparison would express changes as a proportion of total scheduled instances per activity type (i.e., a change rate), which would allow activity types to be compared on a like-for-like basis; this denominator was not available at the time of writing and is noted as a limitation of the current analysis.

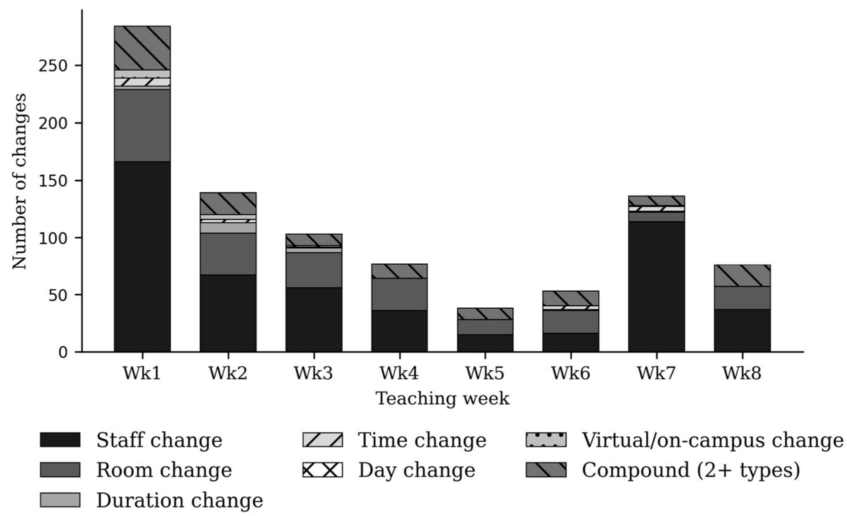

Fig. 1. Weekly distribution of timetable changes by change type, teaching weeks 1–8. Staff and room changes dominate volume throughout, while compound changes (combining two or more change types) and day/time changes, though far less frequent, are concentrated in the earliest teaching weeks.

A secondary pattern visible in Fig. 1 is the spike in staff changes in Week 7, which is comparable in scale to the Week 1 peak; this coincides with a point in the semester where staff cover arrangements are commonly renegotiated and is noted here as a candidate explanation rather than a confirmed cause. These patterns reinforce the importance of temporal context, since a substantial proportion of disruption occurs during the earliest teaching weeks, where changes are likely to have a disproportionate practical impact.

4 Disruptiveness-Weighted Minimal Perturbation

4.1 Limitations of Standard MPP

The standard MPP formulation minimizes the number of changes required to restore feasibility, treating changes primarily through structural measures of deviation, without explicit consideration of their substantially different behavioral impact. Institutional data and feedback from students indicate that this assumption obscures significant differences in practical impact relating to timing, change type and scale.

One source of this differential impact is that many students hold commitments external to their studies — part-time employment, caring responsibilities, or commuting arrangements — that are scheduled around the published day and time of an activity. A change to the day or start time of a session therefore risks conflicting with a commitment the student cannot easily move, whereas a change of room alone leaves the student's wider schedule intact and typically requires only that they walk to a different location. This asymmetry is reflected in the type weighting T, where day and time changes are assigned higher base weights than room changes alone (Sect. 4.3).

4.2 The Disruptiveness Scoring Method

A disruptiveness score D is defined for each timetable change as the product of week weighting, change type weighting, number of students affected and number of weeks affected. Earlier weeks and more compound change types are assigned higher weights. By weighting changes more heavily in earlier weeks and assigning higher scores to more disruptive change types (e.g., day and time changes), the model captures both the temporal and behavioral dimensions of disruption. This approach extends the Minimal Perturbation Problem (MPP) in the timetabling literature by incorporating student-centered metrics, offering a more holistic view of what constitutes a "minimally disruptive" change. This formulation can only be meaningfully applied to post-publication changes, where a student was aware of the previous timetable, since disruption is defined relative to a prior, known state.

A clarification is required regarding the temporal reference point used to determine W. In this model, the week weighting is fixed at the point a change is issued to students: it reflects how early in the teaching period a given disruption lands, independent of how many prior changes that same activity may already have undergone. Where an activity is changed more than once, each change is scored independently at its own issue-week, rather than being measured cumulatively against the originally published timetable. This choice keeps the model usable in an operational setting, since schedules must be updated promptly and a single, fixed point of comparison — the original publication — would assign identical, undiminishing penalties to every subsequent change regardless of when in the term it occurred. The trade-off is that the current formulation does not yet capture the cumulative burden on a student who experiences several smaller changes to the same activity across consecutive weeks; this longitudinal dimension is identified as a direction for future work in Sect. 5. This ensures that temporal sensitivity reflects when disruption is experienced by students, rather than distance from a fixed baseline timetable. This distinction implies that candidate timetable repairs that are equivalent under perturbation-count metrics may be ranked differently when evaluated using a disruptiveness-weighted approach.

4.3 Disruptiveness Score Formula

The disruptiveness score is defined as follows, expressed as the product of four component weights:

$$D = W \times T \times S \times N \tag{1}$$

that is, disruption equals the week weighting multiplied by the type-of-change weighting, the number of students affected, and the number of weeks the change persists, where:

W — the week weighting, decreasing as the teaching semester progresses (W = 8 if the change is issued in Week 1, falling by 1 for each subsequent week, down to W = 1 in Week 8);

T — the type-of-change weighting, derived from the base weights of the individual change types involved (see below);

S — the number of students affected by the change;

N — the number of weeks the change remains in effect.

The total disruptiveness of a proposed set of changes δ is then the sum of the individual disruptiveness scores:

$$min\delta \ \Sigma Di \ = \ \Sigma \,(Wi \times Ti \times Si \times Ni) \tag{2}$$

where δ is a proposed set of timetable changes, Di is the disruptiveness score of change i, and the summation runs over all changes i = 1, …, n in the set.

This formula captures temporal sensitivity through the week of the change, qualitative severity through the type weighting, scale of impact through the number of students affected, and the duration of impact through the number of weeks affected. For compound changes, the type weighting T is computed as the product of the individual change-type weights and is therefore multiplicative rather than additive.

The four component weights (W, T, S and N) are combined multiplicatively rather than additively so that a change which is severe on every dimension — early, compound, and affecting a large cohort — is scored disproportionately higher than a change that is merely severe on one dimension, consistent with the practitioner intuition that such changes are qualitatively, not just additively, worse. The base type weights underlying T in the current model (Day change = 4.97, Time change = 4.46, Duration change = 4.46, Room change = 3.07, Staff change = 2.13, Virtual/on-campus change = 2.07) were assigned by the author based on practitioner judgement and informal student feedback gathered during the study, rather than derived through formal statistical calibration against survey data; they are presented as an illustrative starting scheme rather than a validated weighting, and institutions adopting the model would be expected to calibrate these values against their own student feedback. A consequence of the multiplicative structure is that no individual weight may be set to zero, since this would zero the disruptiveness score regardless of the severity of the other components; all weights in the current scheme are therefore bounded below by approximately 2.

4.4 Patterns of Disruption Across Units

The disruptiveness score was computed for every recorded change in the dataset. Scores were highly right-skewed: across the 50 highest-scoring changes, scores ranged from 4,516 to 77,645, with a median of 7,738 and a mean of 12,506 — indicating that a small number of changes contributed disproportionately to total disruptiveness relative to the bulk of the distribution. The single highest-scoring change in the dataset was a combined room, time and day change to a first-year tutorial in Week 2, affecting 163 students (disruptiveness score: 77,645). By contrast, a Peer Assisted Learning session affected in the same week by only a room and time change, but reaching a substantially larger cohort of 560 students, scored lower overall (53,673) — illustrating that compound change type can outweigh raw student reach in the disruptiveness score.

The application of the disruptiveness score to institutional data revealed stark differences in how disruption is distributed across the curriculum, and demonstrated that frequency of change does not equal total disruptiveness. Some units in the dataset emerged as particularly volatile, with scores exceeding 6,000 points (within the upper quartile of observed changes in the dataset) — placing them among the top quartile of recorded changes by disruptiveness. These units experienced multiple compound changes (combinations of day, room, and time) in the early weeks — the types of changes that students consistently ranked as most disruptive. This demonstrates that changes which are similar in frequency may differ substantially when evaluated using a disruptiveness-weighted metric.

Interestingly, units with high student enrolments were not always the most disruptive. This suggests that the timing and type of change may be more critical than scale alone: a single day change in Week 1, affecting a modest cohort, may be more behaviorally disruptive than a room change in Week 8 for a lecture affecting hundreds of students. This insight challenges traditional assumptions in timetabling optimization, which often prioritize minimizing the number of affected students without considering the wider behavioral and contextual factors of the change — a gap the proposed model is intended to address.

4.5 Implications for Timetabling Practice and Policy

The findings of this study have clear implications for institutional policy and practice. Primarily, universities should prioritize timetable stability in the first four weeks of teaching to reduce disruption. This could involve introducing a change freeze for timetables during this period, improving forecasting for student enrolments, and adopting better staff workload-planning models. Where changes are unavoidable, institutions should aim to minimize behavioral impact by avoiding highly disruptive change types (especially compound day-and-time changes), providing adequate notice, and offering online alternatives where reasonably practical. These implications arise directly from the observed concentration of high-disruptiveness changes in early teaching weeks and from the asymmetry between change types.

5 Conclusion

This paper extends the Minimal Perturbation Problem by introducing a disruptiveness-weighted model for post-publication university timetabling. Empirical analysis demonstrates that changes which are operationally similar can differ substantially in behavioral impact. By quantifying this impact, the MDP model provides a practical decision-support metric that complements existing optimization approaches and supports more human-centric scheduling decisions.

In a heuristic-based timetabling system, the MDP model would be most effectively deployed within the post-publication scheduling window, during the candidate evaluation and selection phase of the heuristic loop — specifically when deciding which change to make among several feasible options.

Moreover, the scoring model could be extended to include additional variables, such as notice period provided, activity type, and cumulative exposure to changes (staff or student). For example, a student who experiences minor changes in three different weeks may incur more cumulative disruption than a student who experiences a single major change. The results suggest that incorporating temporal sensitivity into perturbation evaluation offers a practical means of distinguishing between structurally minimal and behaviorally minimal disruption in dynamic university timetabling. In doing so, the model highlights a limitation of perturbation-based approaches: that proximity to an existing solution is not necessarily a reliable proxy for minimal disruption in practice.

References

1. Abdelhalim, E.A., El Khayat, G.A.: A utilization-based genetic algorithm for solving the university timetabling problem (UGA). Alexandria Engineering Journal 55(2), 1395–1409 (2016).
2. Agarwal, A., McGuire, J., Webb, J.: Understanding the impact of timetable scheduling on university students' academic performance. In: 2023 STARS Conference, Brisbane (2023).
3. Banbara, M., Inoue, K., Kaufmann, B., Okimoto, T., Schaub, T., Soh, T., Tamura, N., Wanko, P.: teaspoon: solving the curriculum-based course timetabling problem with answer set programming. Annals of Operations Research 275(1), 3–37 (2019).
4. El Sakkout, H., Richards, T., Wallace, M.: Minimal perturbation in dynamic scheduling. In: Proceedings of the 13th European Conference on Artificial Intelligence (ECAI 1998), pp. 504–508. John Wiley & Sons, Brighton, UK (1998).
5. Kingston, J.H.: Specifying and solving minimal perturbation problems in timetabling. In: Proceedings of the 11th International Conference on the Practice and Theory of Automated Timetabling (PATAT 2016), pp. 1–14 (2016).
6. Lemos, A., Monteiro, P.T., Lynce, I.: Minimal perturbation in university timetabling with maximum satisfiability. In: International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research, pp. 317–333. Springer, Cham (2020).
7. Lindahl, M., Mason, A.J., Stidsen, T., Sørensen, M.: A strategic view of university timetabling. European Journal of Operational Research 266(1), 35–45 (2018).
8. Müller, T., Rudová, H., Barták, R.: Minimal perturbation problem in course timetabling. In: Burke, E.K., Trick, M.A. (eds.) Practice and Theory of Automated Timetabling V. LNCS, vol. 3616, pp. 126–146. Springer, Heidelberg (2005).
9. Phillips, A.E., Walker, C.G., Ehrgott, M., Ryan, D.M.: Integer programming for minimal perturbation problems in university course timetabling. Annals of Operations Research 252(2), 283–304 (2017).