UCL CENTRE FOR ADVANCED SPATIAL ANALYSIS

UCL

WORKING PAPERS SERIES

Paper 219 - April 20

Optimal Land Use Allocation for the Heathrow Opportunity Area Using Multi-Objective Linear Programming

ISSN 1467-1298

casa

Centre for Advanced Spatial Analysis University College London Gower St London WC1E 6BT
Tel: +44 (0)20 3108 3902 casa@ucl.ac.uk www.casa.ucl.ac.uk

Optimal Land Use Allocation for the Heathrow Opportunity Area Using Multi-Objective Linear Programming

Thomas P. Oléron-Evans[a], Melda Salhab[a,*]

[a] *Centre for Advanced Spatial Analysis, UCL Bartlett Faculty of the Built Environment, 90 Tottenham Court Road, London W1T 4TJ, UK*

Abstract

The London Plan, the Greater London Authority's spatial development strategy for London, has defined Heathrow as an Opportunity Area (OA) – an area with the capacity to support additional homes and jobs – since 2004, but progress on developing the area has been minimal. Uncertainty around the expansion of Heathrow Airport appears to have adversely affected progress. Nevertheless, despite this stagnation, the most recent London Plan stipulates that the Heathrow OA should accommodate 13,000 new homes and 11,000 new jobs.

Using multi-objective linear programming (MOLP), we investigate whether the above figures are achievable given constraints on land availability and land use mix. We further explore how land uses might best be assigned to maximise home, job and GVA creation within the Heathrow OA.

We find that, given 700 ha of available land, as indicated in the London Plan, home and job creation figures can be met. However, a GIS-based investigation of the Heathrow OA reveals that there is insufficient brownfield land to meet these targets, and that development on Green Belt land would very likely be necessary.

Through considering a wide range of different weightings of our objectives and varying the degree to which land use mix is restricted, we present a comprehensive picture of how land can be assigned within the Heathrow OA to maximise key economic objectives. Strong land use allocations from this perspective tend to more heavily feature financial and professional services, other office-based businesses, and shops.

Given that real world land use planning scenarios will always be dependent on factors that cannot be fully captured by a mathematical model, we use MOLP to generate a range of possible allocations reflecting different planning priorities, rather than offering a single "solution". We present our results using a wide range of visualisations to illustrate key trade-offs between different goals, with the aim of promoting MOLP to planners as a valuable tool to support spatial decisions and policy making.

Keywords: Urban Planning, Land Use, Heathrow Airport, Linear Programming, Optimisation, Green Belt

1. Introduction

1.1. Context and Research Questions

In 2004, the Greater London Authority (GLA) published the first iteration of The London Plan, a "Spatial Development Strategy for Greater London", with revised and updated versions following at intervals over the following years (GLA, 2004, 2008, 2011, 2016b, 2017). In order to address the objective of "[accommodating] London's growth within its boundaries without encroaching on open spaces" (GLA, 2004, p. 6), one key policy presented in the original plan was the designation of an initial twenty-eight "Opportunity Areas" (OAs) across Greater London, where the construction of new homes and creation of jobs could be concentrated. These areas were identified on the basis of their capacity for development, their overlap with existing brownfield sites and their accessibility (or potential accessibility) to public transport (GLA, 2004, pp. 39-41).

One of the Opportunity Areas identified in the 2004 plan was "Heathrow/Feltham/Bedfont Lakes", for which initial targets of 5500 new jobs and 930 new homes were set, to be realised by 2016 (GLA, 2004, p. 260). However, while Opportunity Area Planning Frameworks (OAPFs) were produced and adopted for other OAs, by the time of publication of the most recent version of the London Plan (GLA, 2016b), no such framework had been adopted for the Heathrow area. The consultation draft of the next edition of the London Plan appears to confirm that planning in the Heathrow OA has been adversely affected by long-running uncertainty over the proposed expansion of Heathrow Airport, stating that "the area's potential contribution to London's growth [will be reviewed and clarified] when expansion proposals and their spatial and environmental implications are clearer" (GLA, 2017, p. 51).

The latest employment and housing guidelines for the Heathrow Opportunity Area suggest that it should sup-

*Corresponding author. Tel.: +44 (0)20 3108 3876
Email addresses: thomas.evans.11@ucl.ac.uk (Thomas P. Oléron-Evans), melda.salhab.14@ucl.ac.uk (Melda Salhab)

port 13000 new homes and 11000 new jobs (GLA, 2018, p. 26) over a stated area of 700 ha (GLA, 2016b, p. 361), though the shape file of the Heathrow OA (GLA, 2019b) appears to cover a significantly larger area than this. However, all editions of the London Plan have indicated that development should not intrude upon the Green Belt (see, for example, GLA, 2017, p. 62) and the latest shape files of England's designated Green Belt land (MHCLG, 2018) indicate that a significant proportion of the Heathrow OA is in this category. Moreover, shape files of the remaining brownfield sites in London are also available (LPA, 2018) and the amount of such land in the Heathrow OA appears to be very limited.

In this paper, we therefore seek to answer the following research questions:

1. Are the latest guideline figures for homes and jobs in the Heathrow OA achievable and to what extent could development be concentrated on brownfield sites, avoiding development on Green Belt land?

2. How might optimal allocations of land use in the Heathrow OA be specified in order to maximise home and job creation, while also maximising the gross value added (GVA) supported by the available land?

To address these questions, we combine the land use categories and employment density statistics from the London Employment Sites Database (CAG Consultants, 2016) with calculations of GVA per workforce job produced by GLA Economics (2015a), alongside other data sources. We then apply optimisation techniques based on linear programming to produce a range of 'optimal' land use allocation plans for the Heathrow OA.

The potential of linear programming (LP) as a tool for tackling questions of optimal land use planning was first identified in the nineteen-sixties (Schlager, 1965). Since then, it has been applied in a wide range of geographic locations in both urban and agricultural contexts. The technique has both spatially non-explicit versions, which identify optimal proportions of an area to allocate between a number of different land uses (Chuvieco, 1993; Delgado-Matas and Pukkala, 2014; Monaco et al., 2016) and spatially explicit versions, which specify the precise land parcels to be allocated to each land use (Day, 1973; Hayton, 1981; Kumar et al., 2016).

Basic linear programs seek to maximise a single objective function, but some authors have devised variants of the technique that can handle multiple objectives (Glover and Martinson, 1987; Makowski et al., 2000; Aerts and Stewart, 2003). The research questions covered in this paper intrinsically involve multiple objectives, since they seek land use allocation plans that simultaneously excel across a number of different criteria: job and home creation, Green Belt and brownfield development, and GVA growth. We therefore combine and adapt existing multi-objective optimisation techniques, presenting a wide range of visualisations of the possible trade-offs required, in order to provide sophisticated, quantitative and comprehensive answers to these planning questions, allowing policymakers to compare the possible 'optimal' land use allocations that arise from differing prioritisations of their objectives.

1.2. Structure

This paper aims to balance both the practical applications of our methods to the Heathrow context, as well as the presentation of a novel approach to land use multiple-objective linear programming. For content and analysis relating to the Heathrow context, see *Sections 2.1, 3, 4, 6, and 8.* For details regarding our approach to land use multiple-objective linear programming, see *Sections 2.2, 5 and 7.*

The paper is structured as follows:

Section 1: Presentation of context, motivation and research questions; followed by a summary of the paper structure.

Section 2: Comprehensive review of sources relating to the London Plan, the Heathrow OA, and the use of linear programming for land use allocation.

Section 3: GIS-based investigation of the Heathrow OA with regard to Green Belt land, brownfield land, and demographic and economic data.

Section 4: Description of the land use model, collating, calculating and estimating key values for use in later analysis.

Section 5: Specification of: (a) a general multi-objective linear program (MOLP); b) a MOLP for land use allocation; and c) a MOLP for proportional land use allocation.

Section 6: Initial application of MOLP to the research questions, with a discussion of key findings.

Section 7: Specification of a particular approach for solving MOLP, after Glover and Martinson (1987).

Section 8: Application of the Glover-Martinson MOLP to the Heathrow OA, with a discussion of key findings.

Section 9: Summary of answers to the research questions, with discussion of research limitations and future research directions.

2. Context and Literature

2.1. The London Plan and the Heathrow Opportunity Area

2.1.1. The London Plan

The London Plan is the key spatial development strategy for Greater London. It serves as an overarching planning framework that combines the spatial aspects of the Mayor's economic, social, environmental and transport strategies. It aims to ensure consistency across all relevant planning documents and compliance with regional, national, and international policies (GLA, 2016b).

The first London Plan was published in 2004, and due to legislation mandating that the Mayor consistently keeps the plan under review, several versions have been released over the past 16 years (GLA, 2004, 2019c). Most of these releases have been minor updates to an existing plan, but new plans (also known as replacement plans) are also developed, seemingly following each new Mayor taking office. At the time of writing, the current London Plan was published in 2016, but this was an updated version of the London Plan 2011, consolidated with alterations from the interim period (GLA, 2016b). Existing concurrently to the current London Plan is a replacement Plan that the new Mayor is developing. The latest draft was released in July 2019, and upon formal publication, it will replace all previous versions (GLA, 2019c).

Although the stipulation that the London Plan should be kept under review helps ensure that policies remain current, the practice of continuously editing the document can result in inconsistencies. In the context of this paper, the policy targets relating to our area of study have changed significantly across the different iterations of the London Plan. We have chosen to use the most recent figures in our analysis (GLA, 2004, 2016b, 2019c).

The London Plan has an outlook of 20-25 years, but also takes into account longer-term goals, such as London becoming a zero-carbon city by 2050. However, housing targets, which are especially relevant to this study, are presented for delivery within 10 years, to meet London's urgent needs (GLA, 2019c).

2.1.2. Opportunity Areas

All iterations of the London Plan have provided a list of "Opportunity Areas" (OAs), sites that can accommodate at least 2,500 new homes, 5,000 new jobs, or a mix of the two (GLA, 2004, 2016b, 2019c). While these homes and jobs targets have remained a consistent part of the definition of an OA across all versions of the London Plan, other components of the wording have changed. Notably, brownfield land was explicitly mentioned in the definition of OAs in all London Plans until the new London Plan draft, published in Dec 2017, which broadly defined OAs as "the capital's most significant locations with development capacity" (GLA, 2017). The previous version, the 2016 London Plan, defined OAs as "the capital's major reservoir of brownfield land" (GLA, 2016b). This distinction is significant because, as we will establish, based on our calculations for the Heathrow OA, there is not enough brownfield land within the area to accommodate the guideline figures stipulated in the London Plan (GLA, 2016b, 2017).

While the London Plan publishes the list of OAs, it takes only minimal steps towards operationalisation, the burden of which seems to fall upon local boroughs (GLA, 2019c). This is complicated by the fact that OA boundaries can spread across several boroughs, as in the case of the Heathrow OA, which intersects both Hillingdon and Hounslow. Developing an Opportunity Area Planning Framework (OAPF) is typically regarded as the first step towards implementation of an OA. This document is also used as a performance indicator of OA development progress, with each OA being assigned an OAPF number from 1 to 5. A score of 1 represents "nascent", where development potential has been identified and infrastructure options appraisal is underway. An OAPF score of 5 represents "mature" where development potential has been fulfilled and infrastructure is operational (GLA, 2017). Despite the fact that the Heathrow OA was first defined in the 2004 London Plan, it currently has an OAPF progress score of 1, indicating that a planning framework has not yet been developed (GLA, 2019c). We discuss potential reasons for this delay in proceeding sections.

The OAPF progress score is useful for evaluating whether an OA has had a plan developed, but monitoring of plan progress is reported in the London Plan Annual Monitoring Report (AMR) (GLA, 2019a). Data covering net housing completions and pipeline figures are provided, but, interestingly, no information is reported on job creation within OAs. This is a clear limitation, which may have been recognised by the GLA, as the latest AMR 2019a mentions that an OA status report is being developed (GLA, 2019a). Another limitation is that only OAs that have adopted a planning framework are monitored in the AMR. There is no discussion of OAs that do not have a OAPF, so it is difficult to understand the issues causing delays in OAPF development and adoption.

According to the London Plan, the fulfilment of OAPFs should also take into account other key goals, such as protecting the Green Belt (GLA, 2019c). However, certain OAs, such as Heathrow, have a significant amount of Green Belt land within their boundaries. It is unclear to what extent the prevalence of Green Belt land and other potentially development-limiting factors were taken into account when setting homes and jobs guideline figures. In our model, we consider the amount of Green Belt land within the OA boundaries, as part of our effort to evaluate the feasibility of the Heathrow OA targets.

2.1.3. Heathrow Opportunity Area

The first London Plan (2004) identified "Heathrow / Feltham / Bedfont Lakes / Hounslow Town Centre" as an Opportunity Area with indicative estimates of $5,500$ new jobs and 930 new homes over an area of 91 ha, within a 2-year delivery time frame (GLA, 2004). This stands in stark contrast to the 2016 London Plan, which presented the Opportunity Area as "Heathrow" with an indicative employment capacity of $12,000$ and a minimum homes target of $9,000$ over a 700 ha area (GLA, 2016b). The current draft new London Plan identifies the Heathrow OA within the "Heathrow/Elizabeth Line West" strategic development area and the latest indicative figures are $11,000$ jobs and $13,000$ homes. An updated figure for land available for development has not been provided (GLA, 2019c).

The basis upon which these figures have been updated over the years is not discussed in any detail. Regarding the total land available, we could not find any information clarifying the increase from 91 ha to 700 ha. Similarly, the significant increase in the homes and jobs targets is not well explained. The latest draft new London Plans cites the 2017 Strategic Housing Land Availability Assessment (SHLAA) for the homes guideline figure. The jobs figure is based on the latest London Employment Sites (LESD) Database. Given that SHLAA and LESD are separate studies, evaluating the development potential of land for different priorities, it is unclear whether the same plots of land are being considered. In other words, the question of whether 700 ha within the Heathrow OA could genuinely provide both 11,000 jobs and 13,000 homes is of interest, and is addressed by our first research question.

As was mentioned previously, Heathrow has been identified as an OA since 2004, but its OAPF score remains 1, indicating that no plan has yet been approved. The controversy and uncertainty surrounding the expansion of Heathrow Airport has been cited as the critical cause of the delay in the adoption of a planning framework for the Heathrow OA (GLA, 2016b, 2017). In the draft new London Plan, the Mayor emphasises that he will only support expansion plans that do not cause additional noise or a decline in air quality (GLA, 2019c, p. 43, para. 2.1.62).

2.1.4. Heathrow Airport expansion

The potential expansion of Heathrow Airport, with specific regard to the addition of a third runway, has been discussed since at least 2005, when it was suggested in the Department for Transport's Annual Report (Department for Transport, 2005). The series of resulting proposals became the subject of divisive debate that was ultimately politicised, with the government cancelling all plans in 2010, then reintroducing the policy in 2016 (BBC, 2010; Department for Transport and Grayling, 2016). Most recently, in June 2018, the House of Commons voted in favour of the third runway (BBC, 2018).

Proponents of the expansion emphasise that Heathrow Airport is an integral part of Britain's economy, not only in terms of its function as a transport hub, but also as one of the single largest employers in the country. It directly employs 76,000 and further supports 38,000 jobs off-site (Heathrow, 2018, 2019). According to Heathrow Airport's publication supporting the expansion, up to 35,600 new jobs could be created within the airport, and up 22,600 new jobs in the surrounding area (Heathrow, 2018). While the expansion would undoubtedly create employment opportunities, no information is provided explaining the data or methods used to obtain these figures.

As a transport hub, 25% of all UK exports, by value, travel through the airport. Heathrow, which is the busiest two-runway airport in the world, has been operating at capacity for over a decade (Heathrow, 2019). The Department for Transport estimates that all London airports will be full by the mid-2030s (Department for Transport, 2018). Reaching maximum capacity could limit growth and prevent increased connectivity. Operating at capacity would also make it more difficult to manage disruptions and delays, directly impacting passengers and airlines. Consumers could also face increased fares as rising air travel demand is confronted with limited supply (Department for Transport, 2018). Estimates show that the new runway would allow for 265,000 additional flight movements per year, a 56% increase, to reach a new capacity of 740,000 annual flight movements (Heathrow, 2018).

However, increased capacity and cheaper fares would incentivise the use of air travel as a preferred mode of transport, increasing the already high carbon emissions caused by London's airports. Heathrow Airport is currently the UK's largest source of CO_2 emissions (Aviation Environment Federation, 2019). According to the Committee on Climate Change, in the UK, aviation emissions have doubled while emissions from the entire economy have decreased by approximately 40%, over the past two decades (Committee on Climate Change, 2019).

2.1.5. Green Belt land around Heathrow Airport

The current Heathrow expansion plans seem to suggest that development on Green Belt land is inevitable. The published report, "Heathrow Airport Expansion – Our Emerging Plans" (Heathrow, 2018), includes green infrastructure in its 16-component framework, but unlike all the other components, a relevant chapter is not provided. The Green Belt is covered very briefly in the third chapter on planning policy context, but discussion is limited and high-level. The report acknowledges the abundance of Green Belt land around Heathrow, and the site's proximity to the Colne Valley Regional Park. It also cites the construction of Terminal 5 as a precedent, where airport growth was deemed sufficiently important to allow development on former Green Belt land (Heathrow, 2018).

Appendix 3 of the report (Heathrow, 2018) details 36 sites proposed for airport supporting and related development, and all these sites overlap with Green Belt land. Additionally, many of the sites are also in protected areas, such as Areas of Environmental Opportunity, Sites of High Archaeological Potential, and the Colne Valley Regional Park (Heathrow, 2018).

Estimates published by The London Wildlife Trust show that the third runway Heathrow expansion will require 906 ha of total land, approximately half of which (432 ha) is currently designated as Green Belt land (LWT, 2018).

In all likelihood then, any expansion of Heathrow Airport will require Green Belt land. To better understand the prevalence of Green Belt land in the area, we supplement the figures provided in the aforementioned documents with our own GIS based calculations of Green Belt land within the Heathrow OA (see Section 3).

2.1.6. Conclusions

Balancing the many priorities of city planning is, without doubt, an arduous task. Considering the London context, key objectives include: a) providing sufficient, affordable housing; b) creating jobs across the city; c) safeguarding green space; d) ensuring appropriate flexibility in policies to allow for long term economic growth; e) reducing carbon emissions; and so on. The London Plan attempts to meet the ever growing needs of the city by providing a framework and an overall strategy. Achieving all objectives in a timely manner is challenging and, ultimately, trade-offs are likely to occur, as in the example of the Heathrow expansion.

Given the importance of Heathrow Airport to London's economy, the environmental significance of the Green Belt, and the urgent housing and employment needs of the city, the Heathrow OA represents a clear example of a case with multiple, likely conflicting, objectives. A solution, or a set of solutions, that best addresses the majority of priorities is necessary to overcome the stagnation regarding the planning process and allow progress, however defined, to proceed.

Working towards a solution that incorporates numerous priorities is, fundamentally, a multiple objective optimisation problem. Creating a model that quantifies and fully encapsulates all aspects of the decision making problem is unrealistic as human opinions are varied and stakeholder negotiation is essential. However, creating tools that conduct analyses of various portions of the planning problem could provide significant support to policy makers and planners. Furthermore, presenting the results of such analyses through clear and compelling visualisations that highlight critical trade-offs and optimal solutions could allow for nuanced and data-driven discussions.

We seek to apply multi-objective linear programming to find appropriate land use solutions for the Heathrow OA. In the following section, we discuss relevant applications of linear programming for land use optimisation, considering both spatially explicit and non-explicit techniques.

2.2. Linear programming for land use modelling

2.2.1. LP models for land use optimisation

Linear Programming (LP) is an optimisation technique in which the goal is to determine the values that should be assigned to one or more *decision variables* in order that a specified *objective function* is minimised or maximised, subject to a number of *constraints*. The objective function and the constraints, which may be equalities or inequalities, must be linear in the decision variables.

Schlager (1965) was the first author to suggest that land use planning problems could be tackled using LP. Discussing design problems in terms of finding a 'fit' between context and form (after Alexander, 1964), Schlager highlights the complex interactions between the many requirements and restrictions of a typical planning problem as an obstacle to successful design. He then sets out a loose framework for an LP model to address this issue.

In Schlager's framework, a region of interest is partitioned into discrete zones, each of which may be further subdivided to support a number of discrete land uses. Costs per unit area are established for the transformation between specific land uses within each zone, with the objective being to minimise total cost. Constraints are divided into "intra-zonal" requirements, which establish the permitted land uses within a particular zone (whether related to geographical limitations or to planning requirements), and "inter-zonal" requirements, which capture how different land uses interact across the region, along with requirements that specified demands for each land use must be met.

Although Schlager's initial concept had some basic spatial dependency through the consideration of discrete zones, many subsequent authors have successfully applied LP models that are essentially non-spatial. Such applications typically take a similar approach to Schlager, but without the division of a region into separate zones. In this context, given a discrete set of land uses, each decision variable represents the total area of land transformed from one of these uses to another, with the objective function representing the total cost of this transformation (to be minimised), the profitability or employment supported by the land (to be maximised), or similar. Constraints typically relate to the amount of land of different types suitable for transformation, to legal, practical or policy-based restrictions, or to budgetary limitations. The output of such models is a plan indicating the amount of land to be allocated to each use, with no specific indication of how the land uses should be distributed geographically.

Examples of this approach include Chuvieco (1993)'s model for minimising rural unemployment near Castellon, Spain (one of the earlier pieces of work that explicitly used GIS software to enhance LP models), Delgado-Matas and Pukkala (2014)'s model for the optimisation of a traditional land-use system in the Angolan highlands, and Monaco et al. (2016)'s model to optimise agricultural land use in the Milan metropolitan area.

More sophisticated non-spatial LP models have also been applied, incorporating technical enhancements to the basic LP framework. For example, Liu et al. (2007) outline a technique called "inexact chance-constrained linear programming", which combines probabilistic constraints with the concept of "interval linear programming" (in which certain values used in the LP model are not known precisely, but are known only to lie within a certain interval), and apply it to the problem of managing lake areas at the edge of urban conurbations, presenting a case study relating to Wuhan, China. Non-linear programming (NLP) models for land-use management have also been considered. Henseler et al. (2009), for example, use an NLP model designed specifically for the optimisation of agricultural land use to consider the effect of different climate change scenarios on the rural economy in the Upper Danube Basin.

2.2.2. Explicitly spatial land use LP models

Since Schlager's initial work, a wide variety of approaches have been considered for the explicit treatment of space in land use LP models. The most straightforward of these build on Schlager's division of a region into discrete zones or land parcels, incorporating detailed geographical information on each zone. Variables indicating terrain type, elevation and so on, are used to determine the precise form of the LP constraints. Day (1973) applies an approach of this nature to the problem of urban planning on flood plains, in a piece of work that is most notable for its iterative approach. Rather than solving a single problem for the allocation of land use for an entire region, Day tackles a planning problem incrementally, with each step representing a different phase of building work, thus allowing for a more realistic treatment of urban development scenarios.

A slightly different approach to space is taken by Hayton (1981), who applies linear programming to produce land use development plans for the county of Tyne and Wear, UK. Rather than considering continuous space, the author considers 267 discrete sites across the region, with binary decision variables representing the decision to develop or not to develop each site. Such a 0-1 integer programming approach can also be applied on much smaller scales, as demonstrated by Aerts and Stewart (2003). This article considers land use plans for the restoration of part of a former open cast mine in Galicia, Spain. The area of interest is partitioned by means of a square grid, with binary decision variables to represent whether or not each square is allocated to a particular land use. Aerts and Stewart's approach also includes innovative expressions to quantify the compactness of the areas allocated to each land use, allowing for greater control of the patterns of land use distribution that are favoured by the LP model, thus representing a more detailed treatment of the allocation of space.

Recent research has taken an increasingly complex approach to the spatial aspect of land use LP problems. Focussing on the management of urban sprawl, Kumar et al. (2016) formulate an LP model that uses spatial information to consider how well allocated land uses work together, both locally and globally. In common with the approaches discussed above, the model involves the division of the region of interest into a grid of cells, with binary variables to represent whether each potential land use is assigned to each cell. However, Kumar et al. also introduce binary variables to indicate whether pairs of nearby cells are allocated particular pairs of land uses. This approach allows for the identification of solutions that have favourable mixtures of land use over small geographic scales. Such a degree of sophistication comes at a high computational cost, however – the case study presented by the authors (relating to the city of Leander, Texas) involves around four million decision variables and over twelve million constraints – so much of the authors' work is dedicated to methods to make this model more tractable.

2.2.3. Multi-objective land use LP models

The most basic approach for the consideration of multiple objectives in land use LP models is simply to add separate objective functions together to form a compound function, with each of the initial functions multiplied by a weight parameter to indicate their relative importance. However, it can be difficult to choose weights that accurately represent the subjective importance of each goal to a project's stakeholders (even assuming that the interests of those stakeholders were fully aligned), particularly since objectives will generally not be measured in the same units, so must also be standardised somehow to make them comparable. For example, Bammi and Bammi (1979) take this approach in a case study relating to Du Page County, Illinois, simultaneously considering objectives relating to conflict between neighbouring land uses, transportation costs, tax impact, environmental impact and the provision of community facilities. However, the authors also introduce "efficiency constraints", which establish that the ratio of each objective to its minimum feasible value should be roughly equal, thus somewhat reducing the influence of the subjectively chosen weights.

An alternative approach to multi-objective optimisation is to choose one objective function on which to optimise, while integrating others into the LP framework in the form of constraints. These constraints specify a minimum value for each objective function that must be exceeded (supposing that functions are to be maximised) and may be strengthened or relaxed to account for different priorities of the model user. While some authors have expressed misgivings over this blurring of the distinction between objectives and constraints (see Arthur and Nalle, 1997), variations of this approach have nonetheless seen enduring popularity. The aforementioned paper by Hayton (1981) provides a clear example of this approach, while more recent work, such as that of Makowski et al. (2000), is founded on a similar basic concept.

A key issue with the use of LP models to address land use planning problems is their lack of flexibility. A traditional LP model may return a single "optimal" solution, but the value of this solution is limited by how well a complex real world scenario, which inherently involves multiple objectives, can be translated into an LP framework. This limitation was summarised by Downey Brill, Jr. (1979), who criticised contemporary applications of optimisation models to planning problems, "because important planning elements are not captured in the formulations". In the author's view, mathematical optimisation should be used to provide planners with alternative approaches rather than a single answer, complementing rather than replacing the role played by human creativity. Any technique that seems excessively opaque to its intended beneficiaries, and whose solutions are unsuitable owing to criteria that could not be accounted for, is of little value and ultimately risks being ignored and discredited.

For these reasons, various authors have attempted to take a more pragmatic approach to the application of LP, aiming to create methods that fit with the reality of decision making structures and processes. One early piece of work in this vein was produced by MacGregor Smith and Liebman (1978), who describe a method in which expert decision-makers propose a number of potential geographical partitions of a planning area, with a zero-one LP model being applied thereafter to determine which partition should be selected and how activities should be distributed across it. This approach allows for considerable human input into the land use allocation process and ensures that a final land use plan is practically viable, with land uses collected into realistic contiguous patterns, a feature that can be difficult to ensure with pure LP approaches. Nidumolu et al. (2007) discuss an alternative approach, involving the consultation of different stakeholders and the systematic consideration of how the quality of communication between them could affect the adoption of different plans. While their approach raises important issues about the difficulties posed by compromising on objectives that may be of differing importance to different parties, it contributes little to the technical practice of multiple-objective LP, since the authors ultimately reduce their case study (relating to agricultural land management in Andhra Pradesh) to a set of distinct scenarios, each of which is solved as a single objective LP.

Another simple approach to multiple-objective optimisation, which offers at least a degree of flexibility to model users, is to present the complete *Pareto Frontier* of the different objective functions under consideration: the set of feasible points at which no improvement can be made to any one objective function without a deterioration in the value of another. This approach allows model users to see the trade-off between different objectives and potentially therefore to make a more informed choice between solutions. However, while it may seem rational to consider only points on the Pareto Frontier, in practice this may not be the case owing to features of a problem that are not captured in the model framework, as discussed above. Also, while the Pareto Frontier can be clearly visualised for two objective functions (see, for example, Aerts and Stewart, 2003, Fig. 7), for three or more, visualisation is more challenging and less practically useful.

Some more sophisticated multiple-objective LP approaches to land use planning involve the consideration of how close each objective function is to its feasible maximum in percentage terms. A good solution is considered to be one in which all objectives exceed a certain percentage of their maximum, or in which the maximum percentage divergence of any objective from its maximum is minimised. A clear example of the latter approach is provided by Glover and Martinson (1987), who also incorporate weightings for different objective functions, presenting a multiple-objective LP that produces a single solution aiming to balance the priorities of model users. The former approach is more flexible, as demonstrated by Makowski et al. (2000), who present a multiple-objective LP land use planning framework based around the formulation of "nearly optimal solutions". Their innovative concept seeks to characterise and visualise the full range of solutions that come close to attaining the optima of all objectives (in the proportional sense described above), through consideration of *attributes* of the solutions (essentially, sums of groups of decision variables) that would be meaningful to planners. In this way, their method achieves a rare balance between flexibility, clarity and technical elegance, representing a very promising avenue for the practical application of genuine multiple-objective LP to land use planning scenarios.

2.2.4. Conclusions

On consideration of the existing literature, two key issues with the use of multiple-objective LP for land use allocation problems are: a) the need for any method to be flexible enough to accommodate human input, acknowledging that not all aspects of a planning problem can be encapsulated within a mathematical framework; b) the necessity of clear, visual communication of proposed solutions, in terms that are of relevance to planners and other stakeholders.

In the context of our research questions, we adopt two main approaches to address these considerations. Firstly, Pareto Frontiers provide clear visual indications of trade-offs between different objectives and will be a valuable tool to communicate model outputs. Secondly, the method of Glover and Martinson (1987), though designed to produce a single optimal land use allocation, has sufficient flexibility through its incorporation of objective weightings that it allows for a broad investigation of multiple-objective scenarios, such as that of development of the Heathrow OA.

In applying these methods, we will make use of a broad range of visualisations to maximise the clarity of our findings and to explore our research questions from all possible angles.

3. GIS Analysis of the Heathrow OA

3.1. Motivations

Throughout our analysis, we use the 700 ha figure for the amount of available land in the Heathrow OA, which was originally provided in the 2016 London Plan (GLA, 2016b). However, to better address our research questions and to compare the figures published in our references against the geographical realities of the area, we analyse the Heathrow OA using GIS software, with particular focus on the prevalence of Green Belt land and the availability of brownfield sites.

We use the open source software, QGIS, version 2.18, for the GIS analysis (QGIS Development Team, 2019).

3.2. Green Belt and brownfield land

Based on Opportunity Area shapefiles downloaded from the London Datastore (GLA, 2019b), we find that the total area of the Heathrow OA is 7,023 ha. The latest brownfield data available (LPA, 2018), indicates that the brownfield sites within the Heathrow OA boundaries total 41 ha. Most of that land is concentrated in two wards in the Borough of Hillingdon, namely Botwell and Yiewsley. Furthermore, our analysis of the latest Green Belt land data (MHCLG, 2018) reveals that over one third of the land within the Heathrow OA boundaries is Green Belt land, specifically 2,518.5 ha.

The results of our GIS analysis lead us to question the figure of 700 ha of available land, specified in the London Plan. This is significantly higher than the 41 ha of brownfield land that appear to be available, though only about 10% of the total land area of the OA. The high prevalence of Green Belt land in the area is also striking and significant in terms of evaluating its development potential. Nevertheless, absent of further clarifying information, we will evaluate the achievability of the jobs and homes targets for the London Plan's 700 ha figure.

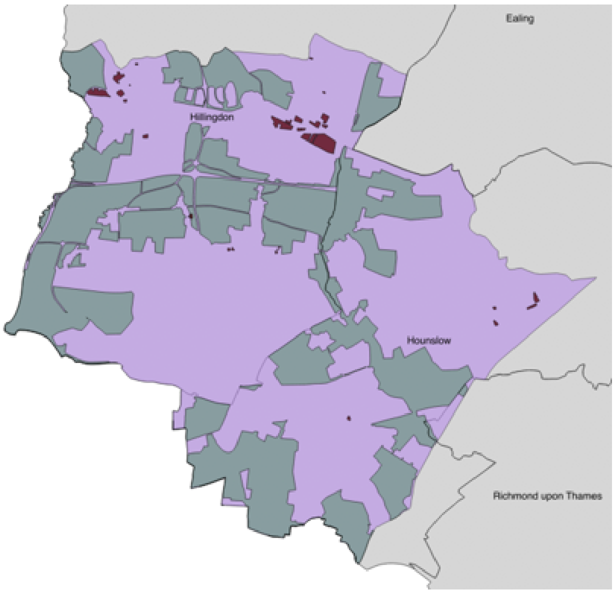

Figure 1: The Heathrow Opportunity Area (purple) with relevant Green Belt (green) and Brownfield Land (brown) indicated

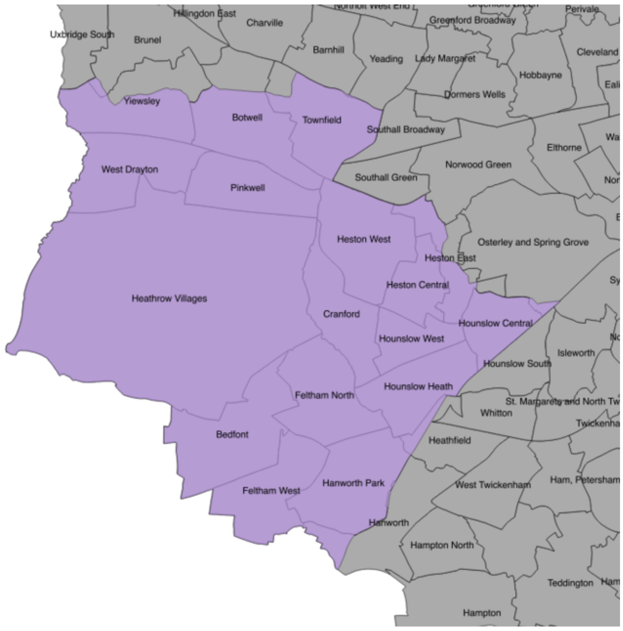

Figure 2: The Heathrow Opportunity Area (purple) and London ward boundaries

Data on wards of Heathrow OA	
Total area	7,250 ha
Total number of jobs (2013)	155,200
Total number of household spaces (2011)	80,552
Average PTAL score (2014)	2.86

Table 1: Overall statistics for 17 wards whose boundaries roughly correspond to the Heathrow OA

3.3. Ward profiles

The extent of the Heathrow OA seems to broadly align with several ward boundaries. The OA overlaps approximately 17 wards, which have a combined area of 7,250 ha, (GLA, 2015b) close our calculation of 7,023 ha for the area of the Heathrow OA. Six of the wards are in the Borough of Hillingdon, including Heathrow Airport, while the remaining eleven wards are in the Borough of Hounslow.

Aligning the OA with the ward boundaries enables us to use ward profiles and atlas data published by the GLA, which provide a range of demographic and economic information (GLA, 2015b). The GLA also provides data segmented into LSOAs (Lower Super Output Area) and MSOAs (Middle Super Output Area), but the boundaries of these areas did not align well with the Heathrow OA.

The table below summarises several relevant indicators:

Taking the values presented in Table 1 as representative of the Heathrow OA, the London Plan targets of creating 11,000 jobs and 13,000 homes would respectively represent increases of 7% and 16% over existing provision.

The London Plan uses a 10-year time frame for its housing targets, which in the case of the Heathrow OA, would translate to 1,300 homes per year. To determine whether that figure is achievable, we analyse the data on annual net additions to dwelling stock for London Boroughs, as the data was not available by ward (MHCLG, 2019). Hillingdon and Hounslow, have respectively seen averages of approximately 780 and 680 new homes per year (net) over the past 10 years. However, in certain years, the number of new homes was far greater, likely due to the completion of large development projects. For example, in 2007-08, Hounslow experienced a net addition of approximately 2,000 new dwellings. While the average annual net additions to dwelling stock across all London boroughs is around 860, certain boroughs have significantly higher rates. Tower Hamlets, for instance, has had an average of about 2,100 new homes per year over the past ten years (MHCLG, 2019). This illustrates that the Heathrow OA target would appear to be achievable on the basis of precedent alone; however, a rate of 1,300 net additions to housing stock per year is well above the norm.

Our analysis also shows that, since the first London Plan was published in 2004, there have been over 23,000 net new homes created in Hillingdon and Hounslow (MHCLG, 2019). Seemingly however, these completions have not been taken into account, since the guideline homes figure for the Heathrow OA has increased over the time period. Nevertheless, since we have established that 1,300 homes per year is broadly feasible on the basis of historic evidence, the proceeding sections will consider whether the figure remains feasible when available land, the jobs target, and other factors are taken into account.

4. The land use model

4.1. Preliminary considerations

Our first research question (see Section 1) asks whether the guideline homes and jobs figures for the Heathrow OA are achievable, while the second explores how land use can be allocated to maximise housing provision, job creation, and the generation of Gross Value Added (GVA).

To answer these questions, it is necessary to establish:

1. Appropriate densities of residential housing for the Heathrow OA;
2. A set of possible land uses for newly developed non-housing land;
3. Estimates for the number of jobs and the GVA that may be supported by each of these land uses, per unit area.

As will be discussed, for the purposes of mathematical modelling and visualisation of results, we will also require:

4. That selected land uses are grouped into broader categories, referred to as "attributes";
5. The identification of a plausible 'default' allocation of non-housing land uses for the Heathrow OA, referred to as a set of "benchmark proportions" for each land use.

The following sections will address these requirements.

4.2. Homes per hectare

To estimate an appropriate number of homes per hectare, we reference the sustainable residential quality density matrix, which provides suggested dwellings per hectare based on public transport accessibility level (PTAL) scores, which measure accessibility to public transport (GLA, 2006). The average PTAL score of the wards that significantly overlap the Heathrow OA is 2.86 (GLA, 2015b). According to the matrix, an appropriate housing density for urban or suburban settings, with a PTAL score between 2 and 3, should be 35-170 dwellings per hectare.

4.3. Non-housing land uses

In 2016, the GLA commissioned CAG Consultants to develop a comprehensive London Employment Sites Database (LESD), bringing together information from a multitude of sources, including city-wide databases and local authorities (CAG Consultants, 2016).

For our analysis, we adopt the land use categories used in this report, condensing the 16 uses listed in its Appendix down to 11, by combining three land uses related to catering into a single category, by combining general and secure residential institutions, and by dropping two other residential categories (since land assigned to housing will be handled separately in our model).

The resulting list of non-housing land uses can be seen in Table 2.

4.4. Jobs per hectare by land use

4.4.1. Data and general principles

The LESD provides employment density estimates (in m^2 per worker) for our selected land uses, across Inner London, Outer London and the Central Activities Zone. We use the employment figures for Outer London, converted into workers per hectare.

When considering the number of jobs supported by a particular unit of land, we are considering the role of land as a base for employment, rather than as a site that could generate work for workers based elsewhere. For example, a block of flats may generate work for plumbers, electricians, nearby schools, doctors and so on, but these workers would be based elsewhere. In particular, this means that land assigned for housing is assumed to directly support no jobs. While there is a risk that this approach may miss certain workers, since it is unclear from the source material exactly how this issue was handled, we adopt this interpretation for simplicity and to avoid double counting.

4.4.2. Hotels

Exceptionally, the employment density estimate for the land use category *LU7: Hotels* is provided in terms of beds per worker, rather than m^2 per worker. In order to convert this to a unit more compatible with our model, we consult the GLA's Housing Standard Report (GLA, 2006), which sets the minimum size of sleeping accommodation for hotels. We consider a standard room of two occupants, which has a minimum size requirement of 10.2 m^2, according to the housing space standards (GLA, 2006). Given that the area of a hotels does not solely comprise bedrooms – other areas such as the reception, restaurants, and service areas must also be housed – we adjust the figure based on the assumption that, on average, approximately 70% of a hotel's floor space is dedicated to bedrooms. This assumption results in a figure of 29.1 m^2 per worker for the *LU7: Hotels* land use category, which we then convert into workers per hectare.

i	Land Use Classifications		Worker/ha	GVA/ha (£M)	Attribute
1	A1	Shops	371.4	14.06	Other Employment
2	A2	Financial and Professional Services	406.3	59.82	Office
3	A3,4,5	Restaurants, Cafes, Bars, Takeaways	371.4	10.53	Other Employment
4	B1	Business – Offices	796.5	46.11	Office
5	B2	General Industrial	180.6	11.71	Manufacturing
6	B8	Storage or Distribution	180.6	9.62	Manufacturing
7	C1	Hotels	223.0	7.25	Other Employment
8	C2	Residential Institutions	144.4	5.26	Other Employment
9	D1	Non-Residential Institutions	144.4	3.56	Other Employment
10	D2	Assembly and Leisure	108.3	3.73	Other Employment
11	SG	Sui Generis	108.3	6.06	Other Employment

Table 2: Workers per hectare and GVA per hectare for eleven selected land uses

4.4.3. Plot ratios

Another necessary consideration is that one hectare of land does not equate to one hectare of work space. Tall office buildings in Central London will have more workers per hectare than office buildings with fewer levels in Outer London. Furthermore, planning laws and other regulations may stipulate the maximum percentage of a plot that may be built up, ensuring a certain degree of green or open space. To account for the different concentrations of work space per unit of land, the LESD provides plot ratios for Central, Inner and Outer London (GLA, 2006, p. 13). We use the Outer London plot ratios to adjust the figures. All land use types have a plot ratio of 65% in Outer London, except *LU4: Business – Offices*, which has a plot ratio of 90%.

4.5. GVA per hectare by land use

The homes and jobs per hectare figures described in the previous section are sufficient to answer the first research question, relating to the achievability of existing targets. To address the objective of maximising GVA, as mentioned in the second research question, we require an estimate of GVA per unit area for each land use.

GVA is a measure of the value of an economy, considering the total goods and services produced. A similar measure is Gross Domestic Product (GDP), which is GVA plus taxes, less subsidies. GVA is often preferred to GDP as a reflection of market conditions because, for example, an increase in tax rate could result in a higher GDP even if there were no corresponding increase in goods and services production (Eurostat, 2019). Essentially, GVA measures all money spent in an economy, which is a key indicator of an economy's health.

To estimate GVA per hectare for each land use, we use data on GVA per workforce job, published by GLA Economics (2015a). The data is provided by section and division, using the categories of the UK Standard Industrial Classification 2007 (SIC07), for London and the UK (GLA, 2015a). We match the categories to the land use classes from the LESD (CAG Consultants, 2016) that were used in the jobs per hectare estimations. We then multiply the GVA per workforce job by the relevant number of workers per hectare to determine the GVA per hectare for each land use.

4.6. Grouping land uses into attributes

We group the eleven non-housing land use types into several categories, or *attributes* (to use the terminology of Makowski et al., 2000, see Section 2.2.3), as seen in Table 2. These attributes will be useful in later sections, for analysing and visualising our results.

The attributes were selected through comparison with a land use map presented in an economic evidence base that was published alongside the current London Plan (GLA, 2016a). The map uses the following employment categories: *Office*, *Manufacturing*, *Other Employment*. We therefore categorise *LU2: Financial and Professional Services* and *LU4: Business – Offices* as *Office*; *LU5: General Industrial* and *LU6: Storage or Distribution* as *Manufacturing*; and all other land uses as *Other Employment*.

4.7. Land use benchmark proportions

While the London Plan does not specify which land uses should be used in the Heathrow OA to meet the guideline figures, land use mix is explicitly encouraged throughout the plan, as in "Policy GG1: Building strong and inclusive communities" (GLA, 2019c, p.14).

Without adaptation to guarantee land use mix, a land use LP will tend to allocate all land to whichever single use is considered to be the most efficient at achieving the objective (or objectives). To avoid this outcome, we will introduce constraints into our LP model to limit how far land use allocation is permitted to diverge from a particular set of *benchmark proportions*.

Specifically, we will define a parameter $\epsilon \in [0, 1]$, the *deviation parameter*, to specify the amount that any land use can deviate from its benchmark proportion. For example, if the *LU7: Hotels* land use were benchmarked at 3% and the deviation parameter were equal to 0.02, only land

i	Land Use Classifications		Benchmark Proportion	Attribute
1	A1	Shops	9.0%	Other Employment
2	A2	Financial and Professional Services	7.4%	Office
3	A3,4,5	Restaurants, Cafes, Bars, Takeaways	7.3%	Other Employment
4	B1	Business – Offices	8.1%	Office
5	B2	General Industrial	5.8%	Manufacturing
6	B8	Storage or Distribution	12.8%	Manufacturing
7	C1	Hotels	2.5%	Other Employment
8	C2	Residential Institutions	18.5%	Other Employment
9	D1	Non-Residential Institutions	3.3%	Other Employment
10	D2	Assembly and Leisure	12.8%	Other Employment
11	SG	Sui Generis	12.5%	Other Employment

Table 3: Benchmark proportions for eleven selected uses, based on the proportion of land allocated to each land use across Greater London (excluding land uses not stated here)

use distributions allocating between 1% and 5% of land to *LU7: Hotels* would be acceptable.

These benchmark proportions, therefore, should represent some form of 'average' or 'typical' land use distribution, which would constitute a plausible baseline land use mix for our model. Of course, what can be considered plausible would vary hugely based on the area of land under consideration; a small plot could plausibly be allocated a single land use, while a whole borough clearly could not.

Absent of more specific information on what could constitute a plausible or typical land use mix for the Heathrow OA, we set the benchmark proportions to represent the current land use mix averaged across the whole of London. We use the latest "Jobs in London by Industry" figures from the ONS (2017), which are presented using the same UK SIC07 categories as the GVA per workforce job data. For each land use, we divide the number of workers by industry by the number of workers per hectare (see Table 2) to obtain the number of hectares allocated to each land use across the city, then convert to percentages.

The resulting benchmark proportions are presented in Table 3. They can also be aggregated to produce benchmark proportions for each attribute (see Table 4), though these will not be specifically used in our LP analysis.

Attribute	Benchmark Proportion
Office	15.5%
Manufacturing	18.6%
Other	65.9%

Table 4: Benchmark proportions from Table 2, aggregated by attribute

Clearly, the concept of benchmark proportions and the deviation parameter is a significant simplification of the way that land is allocated in real planning scenarios. However, as an indicative approach, through considering different values of ϵ, this framework will allow us to gain an insight into the types of land use distribution that may be desirable if a reasonable mix of land uses is to be ensured.

5. Linear Programming Models

5.1. A general framework for practical multi-objective LP

5.1.1. Introduction

As discussed in Section 2.2.3, there is no single way to solve an LP with multiple objectives. Rather, different researchers have proposed different techniques for such problems, to address their own motivations, objectives and intended applications.

In this section, we present a general mathematical framework for a multi-objective linear program (labelled **MOLP0**), with the following two aims in mind:

- To be sufficiently general and flexible that it may encompass various approaches proposed by previous authors;
- To be structured in such a way as to accommodate aspects most relevant for practical applications, rather than to maximise mathematical efficiency.

In Section 5.2, the framework defined here will be applied to land use optimisation, and later to the Heathrow OA specifically. However, the formulation presented here is entirely general.

While the framework is derived from very familiar LP concepts, the particular formulation is original.

Note that this framework it is not intended to provide an approach to 'solving' multi-objective LP problems. It is simply a structure under which to define such problems, in order that different solution strategies may then be applied.

5.1.2. The model framework

The general form of our multi-objective LP, **MOLP0**, is given below.

Throughout, symbols in bold type are used to indicate column vectors of corresponding quantities, of an appropriate length. For example, $\mathbf{x} = [x_1, ..., x_D]^\top$.

MOLP0: General Multi-Objective Linear Program

Choose values for D **Decision Variables**:

$$x_1, \ldots, x_D \in \mathbb{R} \tag{1}$$

Given A **Given Quantities**:

$$g_1, \ldots, g_A \in \mathbb{R} \tag{2}$$

and C **Parameters**:

$$\theta_1, \ldots, \theta_C \in \mathbb{R} \tag{3}$$

Also define B **Derived Quantities**:

$$d_1(\mathbf{x}, \mathbf{g}, \boldsymbol{\theta}), \ldots, d_B(\mathbf{x}, \mathbf{g}, \boldsymbol{\theta}) \in \mathbb{R} \tag{4}$$

These quantities are specified functions of the decision variables, the given quantities and/or the parameters.

Specify E **Objective Functions** to be maximised:

$$Z_i(\mathbf{x}) = \sum_{j=1}^{D} z_{ij}(\mathbf{g}, \boldsymbol{\theta}) x_j \ , \ \ \forall i \in \{1, \ldots, E\} \tag{5}$$

where the z_{ij} are specified functions of the given quantities and the parameters.

Subject to G **Inequality Constraints**:

$$\sum_{j=1}^{D} p_{ij}(\mathbf{g}, \boldsymbol{\theta}) x_j \leq s_i \ , \ \ \forall i \in \{1, \ldots, G\} \tag{6}$$

for specified values $s_1, \ldots, s_G \in \mathbb{R}$, and where the p_{ij} are specified functions of the given quantities and the parameters.

and F **Equality Constraints**:

$$\sum_{j=1}^{D} p'_{ij}(\mathbf{g}, \boldsymbol{\theta}) x_j = r_i \ , \ \ \forall i \in \{1, \ldots, F\} \tag{7}$$

for specified values $r_1, \ldots, r_F \in \mathbb{R}$, and where the p'_{ij} are specified functions of the given quantities and the parameters.

And record E **Single Objective Optima**:

$$Z_{1,\max}, \ldots, Z_{E,\max} \in \mathbb{R} \tag{8}$$

These are the optimal values that are attained when this linear program is solved for each objective function separately, where these optima exist and are finite.

5.1.3. Further mathematical notation

For mathematical convenience, it will be helpful to establish some supplementary notation to accompany model **MOLP0**.

Firstly, considering the Single Objective Optima, the set of feasible points at which each optimum is obtained is denoted:

$$\begin{aligned} X_i = &\{\mathbf{x} \in R : Z_i(\mathbf{x}) = Z_{i,\max}\}, \\ &\forall i \in \{1, \ldots, E\} \end{aligned} \tag{9}$$

where R is the feasible region: the region of $\mathbb{R}^D$ in which all equality and inequality constraints are satisfied.

Note that, if the relevant LP is not feasible or is not bounded, then $Z_{i,\max}$ does not exist and the corresponding set X_i is not defined.

If X_i consists of a single point, this point is denoted:

$$\mathbf{x}_{i,\max} = [x_{i,1,\max}, ..., x_{i,D,\max}]^\top \in R \tag{10}$$

5.1.4. Discussion of the general model

Note that, in order to satisfy the aims set out in Section 5.1.1, this definition of a multi-objective linear program contains sections that are mathematically redundant, but which are nonetheless practically important in real scenarios. Specifically, there is no mathematical difference between the status of the "Given Quantities" and the "Parameters" in this linear program, but in real life problems, the given quantities may represent measured values of relevant quantities in the real world, while the parameters would define different ways of approaching the problem, corresponding to different assumptions or priorities. Similarly, there is no mathematical need to explicitly define a category of "Derived Quantities", but in real life problems it can be useful to clearly set out important quantities to be calculated from more fundamental measurements and values.

Unlike linear programs with a single objective function, there is no single method for solving problems with multiple objectives. Finding good feasible solutions will generally require a trade-off between the multiple objectives and, as we have discussed, there are different ways in which such trade-offs can be defined and addressed.

5.1.5. Pareto dominance and the Pareto Frontier

The concepts of Pareto dominance and Pareto Frontiers, as discussed in Section 2.2.3, are formally defined below, in the context of our general multi-objective LP framework:

Definition 5.1. *Consider a multi-objective LP of the form* **MOLP0***, with feasible region R. Given two feasible solutions $\mathbf{x}, \mathbf{y} \in R$, $\mathbf{y}$* ***Pareto dominates*** *$\mathbf{x}$ if and only if:*

$$Z_i(\mathbf{y}) \geq Z_i(\mathbf{x}) \ , \ \ \forall i \in \{1, \ldots, E\}$$

and:

$$\exists j \in \{1, \ldots, E\} \text{ such that } Z_j(\mathbf{y}) > Z_j(\mathbf{x})$$

Definition 5.2. *Consider a multi-objective LP of the form* **MOLP0***, with feasible region R. The* ***Pareto Frontier*** *of the multi-objective LP is defined as:*

$$\mathcal{P} = \{\mathbf{x} \in R : \nexists \mathbf{y} \in R \text{ for which } \mathbf{y} \text{ Pareto dominates } \mathbf{x}\} \quad (11)$$

In other words, to say that a feasible solution $\mathbf{y}$ Pareto dominates a feasible solution $\mathbf{x}$ means that $\mathbf{y}$ is at least as good as $\mathbf{x}$ on all objective measures, and strictly better than $\mathbf{x}$ for at least one of them. The Pareto Frontier is the set of all feasible solutions that are not Pareto dominated by any other feasible solution.

A natural argument could be made that it is not rational to select a feasible solution that does not lie on the Pareto Frontier, since, for such a solution, there must exist a solution on the Frontier that is preferable on at least one objective and no worse on any other. This would suggest that feasible solutions that do not lie on the Frontier can be disregarded from consideration when applying multi-objective LP. However, as discussed in Section 2.2.3, for real life applications, this may not be the case, as there will likely be important features of a problem that are not captured or cannot be captured within a linear programming framework. Such unmodelled features could mean that a Pareto dominated solution does in fact represent a reasonable outcome.

5.2. Applying the general framework for land use optimisation

5.2.1. The land use LP model

In this section, we set out the structure of the particular multi-objective LP that underpins the methodology of this article. This LP, set out according to the structure defined in Section 5.1, seeks to optimise the distribution of land across the eleven non-housing land uses discussed in Section 4.3, along with housing land use, considering objective functions relating to home, job and GVA generation.

As in the general model, symbols in bold type are used to indicate column vectors of corresponding quantities, of an appropriate length.

MOLP1: The Land Use Multi-Objective LP

$N + 1$ **Decision Variables**:

$$x_0, ..., x_N \quad (12)$$

These represent the amount of land (in hectares) allocated to each of $N + 1$ land uses, where land use 0 is residential and the remaining land uses correspond to those in Table 2.

Given Quantities:

- The amount of land available for development (in hectares): $T > 0$;
- The number of jobs supported by one hectare of land allocated to land use i: $j_i \geq 0$;
- The GVA (in £) supported by one hectare of land allocated to each land use i: $v_i \geq 0$;
- The density of homes per hectare on land assigned for housing: h;
- A minimum target of homes to be built: $H_{\min} \geq 0$;
- A minimum target of jobs to be created: $J_{\min} \geq 0$;
- Benchmark proportions for each non-housing land use: $\bar{q}_1, \ldots, \bar{q}_N \in [0, 1]$ (see Section 4.7).

Parameter:
As discussed in Section 4.7, the deviation parameter, $\epsilon \in [0, 1]$, specifies the absolute deviation permitted between the actual proportions, q_i, of non-housing land allocated to each land use and their corresponding benchmark proportions $\bar{q}_i$.

Derived Quantities:

- The number of homes built:

$$H(\mathbf{x}) = hx_0 \quad (13)$$

- The number of jobs created:

$$J(\mathbf{x}) = \mathbf{j} \cdot \mathbf{x} \quad (14)$$

- The total GVA created:

$$V(\mathbf{x}) = \mathbf{v} \cdot \mathbf{x} \quad (15)$$

- The proportion of land allocated to each non-housing land use, disregarding the land allocated to housing:

$$q_i = \frac{x_i}{T - x_0}, \qquad \forall i \in \{1, \ldots, N\} \quad (16)$$

Objective Functions, to be maximised:

1. $Z_H(\mathbf{x}) = H(\mathbf{x})$ (homes built)
2. $Z_J(\mathbf{x}) = J(\mathbf{x})$ (jobs created)
3. $Z_V(\mathbf{x}) = V(\mathbf{x})$ (GVA generated)

Equality Constraints

- All available land must be allocated to one of the selected uses:

$$T = \sum_{i=0}^{N} x_i \quad (17)$$

Inequality Constraints

- Job and housing targets must be respected:

$$J(\mathbf{x}) \geq J_{\min} \quad (18)$$
$$H(\mathbf{x}) \geq H_{\min} \quad (19)$$

- The proportion of land allocated to each land use (excepting housing land use, x_0) must not deviate from the corresponding benchmark proportion by an amount greater than the deviation parameter:

$$\bar{q}_k - \epsilon \leq q_k \ , \qquad \forall i \in \{1, \ldots, M\} \tag{20}$$
$$\bar{q}_k + \epsilon \geq q_k \ , \qquad \forall i \in \{1, \ldots, M\} \tag{21}$$

- Non-negativity constraints:

$$x_i \geq 0 \ , \qquad \forall i \in \{0, \ldots, N\} \tag{22}$$

Single Objective Optima:

We denote the maxima of Objective Functions 1-3 (each obtained by disregarding the other two objectives and solving the resulting single objective linear program) as $H_{\max}$, $J_{\max}$ and $V_{\max}$ respectively.

5.2.2. Further mathematical notation

As in Section 5.1.3, the sets of feasible solutions at which the maxima $H_{\max}$, $J_{\max}$ and $V_{\max}$ are obtained will be denoted: X_H, X_J, X_V. If these sets contain unique points, these solutions will be denoted: $\mathbf{x}_H$, $\mathbf{x}_J$, $\mathbf{x}_V$.

5.2.3. Model quantities to represent the Heathrow OA

Table 5 summarises the values of the key quantities from **MOLP1** that we use to represent the Heathrow OA.

Quantity	Value	Justification
T	700	Section 2.1.3
j_i	various (see Table 2)	Section 4.4
v_i	various (see Table 2)	Section 4.5
h	$35 - 170$	Section 4.2
$H_{\min}$	13000	Section 2.1.3
$J_{\min}$	11000	Section 2.1.3
$\bar{q}_1, \ldots, \bar{q}_N$	various (see Table 3)	Section 4.7

Table 5: Quantities from **MOLP1** and the values used to represent the Heathrow OA

5.2.4. A proportional land use LP model for non-housing land uses

One important observation to be made about the Land Use LP, **MOLP1**, is that there is no direct interaction between land used for housing and non-housing land.

To understand this point, first observe that land assigned for housing (land use 0) supports neither jobs nor GVA, $j_0 = v_0 = 0$, and that land assigned for non-housing use does not support housing. Therefore, the area of land allocated to housing, x_0, does not feature in the objective functions relating to jobs and GVA, Z_J and Z_V, and the areas of land allocated to other uses, $x_1, \ldots, x_N$, do not feature in the objective function relating to housing, Z_H.

Secondly, observe that land assigned for housing is not benchmarked, as other land uses are. This means that, while the proportions of land allocated to non-housing uses are constrained in an interrelated way, through the benchmark inequalities (20) and (21), the only constraint on land assigned for housing is that it should meet its target $H_{\min}$ (inequality (19)).

Next, consider that the proportions q_i of land allocated to each land use (16) are calculated only after land assigned for housing has been deducted from the total land available. This is intended to represent a simplified scenario in which planners first decide what proportion of their land should be allocated to housing and then consider how the remaining land should be used. It means that the benchmark proportions on other land uses do not limit the amount of land that may be used for housing in any way.

Considering these observations, and noting that the only constraint in **MOLP1** that combines both housing and other land uses is the constraint on the total land available (17), we observe that housing and non-housing land uses are effectively completely decoupled. All else being held equal, by the linearity of the problem, if x_0 is fixed at a particular value and and the resulting linear program is solved to maximise Z_J or Z_V, while the optimal values of the variables $x_1, \ldots, x_N$ will naturally depend on the value of x_0, the proportions, $q_1, \ldots, q_N$, will not.

This suggests that the questions of maximising jobs and GVA should be considered separately from the question of maximising housing, and that it is the proportions of non-housing land allocated to each land use (q_i) that are of primary relevance.

MOLP1 can therefore be reframed to consider the relationship between the proportions of non-housing land allocated to each land use and the resulting jobs and GVA created per hectare, by ignoring the housing and jobs targets, disregarding the housing objective function ($H(\mathbf{x})$), and setting the total available land to one hectare ($T = 1$). In this setting, the decision variables represent the proportion of land to be allocated to each land use (and are equal to the q_i), for a region of arbitrary size, and the objective functions $\hat{J}(\mathbf{x})$ and $\hat{V}(\mathbf{x})$ represent the jobs and GVA created per hectare.

This altered version of **MOLP1** will be referred to as the Proportional Multi-Objective Land Use LP (**MOLP1-P**), and will be used throughout the majority of this article to examine the relationship between jobs and GVA, without consideration of housing, in the Heathrow OA.

A full specification follows:

MOLP1-P: Proportional Land Use Multi-Obj. LP

N **Decision Variables**:

$$q_1, \ldots, q_N \tag{23}$$

The proportions of non-housing land allocated to each of N land uses, as in **MOLP1**.

Given Quantities:

Jobs per hectare and GVA per hectare for each land use, j_i and v_i, and the benchmark proportions $\bar{q}_i$, as in **MOLP1**.

Parameters:
The deviation parameter, ϵ, as in **MOLP1**.

Derived Quantities:

- Overall jobs per hectare:

$$\hat{J}(\mathbf{q}) = \mathbf{j} \cdot \mathbf{q} \tag{24}$$

- Overall GVA per hectare:

$$\hat{V}(\mathbf{q}) = \mathbf{v} \cdot \mathbf{q} \tag{25}$$

Objective Functions, to be maximised:

1. $Z_{\hat{J}}(\mathbf{q}) = \hat{J}(\mathbf{q})$ (jobs)
2. $Z_{\hat{V}}(\mathbf{q}) = \hat{V}(\mathbf{q})$ (GVA)

Equality and **Inequality Constraints**:

- Proportional land uses must sum to 1:

$$1 = \sum_{i=1}^{N} q_i \tag{26}$$

- Benchmark inequalities, as in **MOLP1**:

$$\bar{q}_k - \epsilon \leq q_k \ , \qquad \forall i \in \{1, \dots, M\} \tag{27}$$
$$\bar{q}_k + \epsilon \geq q_k \ , \qquad \forall i \in \{1, \dots, M\} \tag{28}$$

- Non-negativity constraints:

$$q_i \geq 0 \ , \qquad \forall i \in \{1, \dots, N\} \tag{29}$$

Single Objective Optima:
We denote the maxima of the objective functions (each obtained by disregarding the other objective and solving the resulting single objective linear program) as $\hat{J}_{\max}$ and $\hat{V}_{\max}$ respectively. The corresponding feasible solutions are denoted as $\mathbf{q}_J$ and $\mathbf{q}_V$ (or the set of corresponding feasible solutions, Q_J and Q_V).

5.3. Conclusions

In this section, we have defined three frameworks for multi-objective LP: a general model, applicable to problems in many domains; a model specific to land use optimisation problems, which can be applied to the Heathrow OA; and an adaptation of the land use model that considers proportional rather than absolute areas and which disregards land assigned for housing. However, we have not yet specified a method to identify solutions (i.e. land use allocation plans) from these frameworks, taking full account of their multiple objectives.

In the next section, we use the Land Use Multi-Objective LP to provide initial answers to our research questions, before undertaking a more detailed analysis in Section 7, where we will apply the method of Glover and Martinson (1987) to explore our second research question in more depth.

6. Initial consideration of research questions

6.1. Research Question 1:

Are the latest guideline figures for homes and jobs in the Heathrow OA achievable and to what extent could development be concentrated on brownfield sites, avoiding development on Green Belt land?

6.1.1. Summary of findings

Our analysis demonstrates that the targets of 13,000 new homes and 11,000 new jobs across 700 ha of land are achievable. Indeed, these targets should be relatively comfortable over an area of this size.

As we will discuss, 90.3 ha represents a theoretical minimum amount of land needed to achieve the figures, though this would be at a fairly high housing density and assumes complete freedom over land use mix. However, assuming lower density housing and very strict limitations on the permitted land use mix yields an upper bound on the amount of land required of 414.9 ha, still well below the available total of 700 ha (see Table 6).

Recall however that, according to the latest data available from the London Planning Authorities (2018), only 41 ha of brownfield land is available within the Heathrow OA (see Section 3.2). Our analysis therefore suggests that it would not be possible to achieve the guideline figures using only brownfield land.

This implies either that significant amounts of brownfield land or otherwise undeveloped land are available for development, but are not recorded in the data; that existing developed land (whether commercial, industrial, agricultural or residential) would need to be repurposed; or that achieving the targets would necessitate building on some of the 2,518.5 ha of Green Belt land in the Heathrow OA.

6.1.2. Minimum land required to achieve homes and jobs targets

The two key factors affecting the amount of land needed to achieve the targets are housing density (h) and permitted deviation from the benchmark proportions, (specified by the deviation parameter, ϵ). Higher housing density would allow the housing target to be met using less land. Similarly, as the model is permitted to deviate further from the benchmarks (increasing ϵ), land may be allocated to land uses that more efficiently generate jobs, so less land would be required to meet the jobs target. Note, however, that the parameter ϵ does not affect the amount of land required for housing.

Table 6 shows the minimum land required to achieve the guideline figures across a range of residential densities and ϵ values. Five values of ϵ (0.00, 0.05, 0.10, 0.20, 1.00) and five values of h (35.00, 68.75, 102.50, 136.25, 170.00) were considered, resulting in a total of twenty-five scenarios. The values of ϵ were chosen after exploratory analysis to represent a broad range of outcomes, focusing mainly on lower values of ϵ, which are likely to be more

Permitted deviation from benchmark proportions (ε)	Maximising:	Residential Density (Housing Units / ha) 35.00	68.75	102.50	136.25	170.00
0.00	Homes	371.4	189.1	126.8	95.4	76.5
	Jobs	43.5	43.5	43.5	43.5	43.5
	TOTAL	414.9	232.5	170.3	138.9	119.9
0.05	Homes	371.4	189.1	126.8	95.4	76.5
	Jobs	33.7	33.7	33.7	33.7	33.7
	TOTAL	405.1	222.8	160.5	129.1	110.1
0.10	Homes	371.4	189.1	126.8	95.4	76.5
	Jobs	27.6	27.6	27.6	27.6	27.6
	TOTAL	399.0	216.7	154.4	123.0	104.1
0.20	Homes	371.4	189.1	126.8	95.4	76.5
	Jobs	22.0	22.0	22.0	22.0	22.0
	TOTAL	393.4	211.1	148.8	117.4	98.4
1.00	Homes	371.4	189.1	126.8	95.4	76.5
	Jobs	13.8	13.8	13.8	13.8	13.8
	TOTAL	385.2	202.9	140.6	109.2	90.3

Table 6: Minimum land required (in hectares) to achieve homes and jobs targets for various housing densities and values of the deviation parameter, ϵ

applicable to real world situations in which a diverse land use mix must be maintained. The values of h were chosen uniformly across the range of realistic values established in Section 4.2.

The minimum land required to achieve the housing target was calculated simply by dividing the target (13,000) by the specified density. The minimum land required to achieve the job and GVA targets was produced by determining the optimal jobs per hectare figures using **MOLP1-P** (see Section 5.2.4) for each value of ϵ, then dividing the target (11,000) by this value.

We observe that in the least residentially dense scenario, where 35 housing units are built per hectare, 371.4 ha are needed to meet the 13,000 new homes target. At the other end of the spectrum, where 170 units are built per hectare, only 76.5 ha would be needed to meet the housing target. 76.5 ha is therefore the minimum land required, given the acceptable housing density range, as specified by the GLA Housing report (GLA, 2006).

To meet the jobs target, considerably less land is needed. Assuming land use allocation at benchmark proportions ($\epsilon = 0$), 43.5 ha of land is required. Permitting deviation from the benchmarks allows the allocation of land uses to be more efficient with respect to job creation. At a 20 percentage point deviation from the benchmarks ($\epsilon = 0.2$), 22 ha would be required to create 11,000 new jobs. When full deviation is permitted ($\epsilon = 1.0$), effectively meaning that all land is allocated to the single most efficient land use for job creation, only 13.8 ha would be required to meet the target.

6.1.3. The Heathrow expansion

The Heathrow third runway expansion will require 906 ha of land (not all of which lies within the Heathrow OA), of which 432 ha is currently Green Belt land (LWT, 2018). The achievability of the guideline figures in the case of expansion therefore depends on whether some or all of the 700 ha earmarked for home and job creation by the London Plan (2016b) will be used for the expansion.

The London Plan (2016b; 2017) only provides the available land figure, 700 ha, and a geometry outlining the entire OA, which includes built-up land, Green Belt land, and brownfield land. Specific sites for home and job creation are not defined, so it is unclear to what extent the 906 ha of land required for the Heathrow expansion overlap with the 700 ha mentioned in the London Plan (2016b; 2017).

6.2. Research Question 2:

How might optimal allocations of land use in the Heathrow OA be specified in order to maximise home and job creation, while also maximising the gross value added (GVA) supported by the available land?

6.2.1. Initial considerations

While Research Question 1 focused on the achievability of the job and housing targets, Research Question 2

Permitted deviation from benchmark proportions (ε)	Maximising:	Residential Density (Housing Units / ha)				
		35.00	68.75	102.50	136.25	170.00
0.00	Homes	23,000	45,100	67,300	89,500	111,600
	Jobs	83,000	129,000	145,000	153,000	158,000
	GVA (£M)	4,800	7,400	8,200	8,900	9,200
0.05	Homes	23,300	45,800	68,300	90,800	113,300
	Jobs	107,000	167,000	187,000	197,000	204,000
	GVA (£M)	6,700	10,400	11,600	12,300	12,600
0.10	Homes	23,500	46,200	68,900	91,600	114,300
	Jobs	131,000	204,000	228,000	241,000	249,000
	GVA (£M)	8,400	13,100	14,700	15,500	16,000
0.20	Homes	23,700	46,600	69,500	92,400	115,300
	Jobs	164,000	256,000	287,000	303,000	312,000
	GVA (£M)	11,600	18,000	20,200	21,300	22,000
1.00	Homes	24,000	47,200	70,300	93,500	116,700
	Jobs	262,000	407,000	457,000	482,000	497,000
	GVA (£M)	19,700	30,600	34,300	36,200	37,300

Table 7: Maximum homes, jobs, and GVA achievable when each is maximised independently, for various housing densities and values of the deviation parameter, ϵ

considers how best to maximise home, job and GVA generation.

The first thing to observe is that the question of maximising housing provision is trivial when considered simultaneously with those of maximising job and GVA generation. As discussed in Section 5.2.4, housing provision and job and GVA provision are essentially decoupled in our optimisation framework, such that it makes sense only to specify a desired amount of housing, allocate an appropriate amount of land, and then decide how to allocate non-housing land uses over the remaining land, considering the objectives of job and GVA generation. Given a particular plan of how to allocate non-housing land uses (the proportions q_i), the trade-off between housing on the one hand and job and GVA generation on the other is entirely linear in x_0 (the area allocated to housing).

The question of how to balance the objectives of job and GVA generation is much broader, since these goals are heavily interrelated in the optimisation framework.

In this section, we consider the maximum values of each objective function ($H_{\max}$, $J_{\max}$ and $V_{\max}$) that can be achieved when the others are disregarded, across a spectrum of different scenarios. We then undertake an initial examination of the complex trade-off between job and GVA generation, before applying a more sophisticated approach, based on the work of Glover and Martinson (1987) in the remainder of the article.

6.2.2. Maximising homes, jobs, and GVA independently

We calculate the maximum homes, jobs, and GVA figures that can be achieved over our area of 700 ha, while meeting the specified targets, for the same twenty-five scenarios set out in Table 6. These values are simply $H_{\max}$, $J_{\max}$ and $V_{\max}$, produced by **MOLP1**, for the appropriate values of ϵ and h.

The results are presented in Table 7. For example, at a housing density of 102.5 homes per hectare, with $\epsilon = 0.1$, the maximum amounts of homes, jobs and GVA that can be created while meeting the housing and jobs targets are 68,900, 228,000, and £14.7 bn respectively. Note that these figures cannot be achieved simultaneously, they are simply the maxima of each objective when others are disregarded. Indeed, while jobs and GVA are naturally somewhat positively correlated, the housing total will be at its minimum value (the target of 13,000) when the maxima for jobs or GVA are achieved and vice versa.

Table 7 shows that, depending on how much freedom is permitted regarding the land use mix and the density of residential housing, the maximum number of homes that can be created on 700 ha of land, while meeting the jobs target, ranges from 23,000 (under the strictest assumptions) to 116,700 (under the loosest assumptions). The maximum number of jobs that can be created while meeting the housing target ranges from 83,000 to 497,000, while the maximum GVA that can be produced ranges from £4.8 bn to £37.3 bn.

Studying the influence of the deviation parameter ϵ

and the housing density separately offers additional insights. Considering the column representing the minimum housing density of 35 in Table 7, we see that the maximum number of homes that can be built increases by just 4%, as ϵ increases from its minimum value of 0 (where all non-housing land is allocated according to the benchmark proportions) to 1 (where there are no restrictions on non-housing land use mix).

This can be understood in the context of our previous discussion of the decoupling of the residential and non-housing optimisations. The value of ϵ does not directly influence the number of homes that can be built. Rather, its influence is indirect, since higher values of ϵ mean that less land is needed to meet the job creation target, leaving more land to be allocated to housing. As would be expected, for the same housing density, the impact of ϵ on the maximum number of jobs and value of GVA that can be created is more pronounced, with both these figures more that trebling as ϵ varies from 0 to 1.

On the other hand, fixing $\epsilon = 0$ and increasing housing density from its minimum to its maximum value results in an almost four-fold increase in housing (a linear increase with h). In the same scenario, job and GVA maxima increase by 90% and 92%, respectively. Although housing density does not interact directly with jobs or GVA, higher densities allow the housing target to be met over a smaller land area, freeing up land for jobs and GVA generation.

6.2.3. Pareto Frontiers for job and GVA generation

We now set aside considerations of housing to focus on the relationship between job and GVA generation on land allocated to non-housing uses. Table 7 has established the maximum housing figures that may be achievable at different residential densities, and it remains to consider how to maximise and balance the jobs and GVA supported by non-housing land on a per hectare basis.

We begin by examining the Pareto Frontiers (see Section 5.1.5) of jobs and GVA per hectare, for $\epsilon = 0.00$, 0.01, 0.02, ..., 1.00. Each Pareto Frontier was determined by applying **MOLP1-P** repeatedly for each value of ϵ, with the jobs per hectare figure restricted to exceed fixed values (by means of an additional constraint), ranging from (a) the maximum jobs per hectare that could be achieved when maximising GVA to (b) the maximum jobs per hectare that could be achieved when GVA was disregarded. For each of these job creation figures, the maximum GVA per hectare (i.e. $\hat{V}_{\max}$) produced by the LP was noted, and these GVA maxima were plotted against the job figures to produce the Frontier.

Figure 3 displays these Pareto Frontiers, demonstrating the relationship between GVA per hectare and jobs per hectare as the deviation parameter increases. Diamond shaped points indicate values of ϵ for which the Pareto Frontier is a single point, indicating that identical or equivalent land use distributions simultaneously maximise both jobs and GVA in the corresponding scenario. Note in particular that $\epsilon = 0$ corresponds to such a point, since proportional land uses are fixed at the benchmarks in this case. However, singleton Pareto Frontiers are also observed for values of ϵ from 0.26 to 0.42.

Though not easy to visualise, Pareto Frontiers for values of ϵ from 0.43 to 1 are collinear, simply expanding in length as ϵ increases. Behaviour in this region will be discussed in greater detail in Section 8.3.2.

From the plot, we observe a strong positive relationship between jobs and GVA per hectare for values of ϵ below 0.43. The Pareto Frontiers in this region are either narrow compared to the absolute values of jobs and GVA per hectare that they cover, or they are singletons, suggesting that any gains that can be made in one of these objectives at the expense of the other are relatively small or non-existent when land use mix is strictly to moderately restricted.

For values of ϵ between 0.1 and 0.25, the Pareto Frontiers appear relatively "flat". Taking account of the scale of the axes, this indicates that, for a particular value of ϵ in this range, even when job creation is maximised, a proportionally high level of GVA is possible (compared to its feasible maximum), while there is a proportionally greater hit to job creation when exclusively maximising GVA.

For $\epsilon \geq 0.43$ the picture is rather different. From this point onwards, there is a clear and increasing trade-off to be considered between the two objectives. The maximum possible GVA and jobs per hectare may be achieved when ϵ is equal to 1 (in fact, Pareto Frontiers are identical for all $\epsilon \geq 0.93$), for which there is no restriction on land use mix. At this point, all land may be allocated to the single, most efficient land use for a particular objective.

The maximum possible GVA per hectare is £59.8M, at which point all land is allocated to *LU2: Financial and Professional Services*. However, this corresponds to only 406 jobs per hectare. Comparatively, when all land is allocated to *LU4: Business – Offices*, we attain the maximum possible jobs per hectare of 796.5, though with an associated GVA per hectare of only £46.1M.

Higher values of the deviation parameter ϵ likely represent scenarios that are not practically relevant for urban planning over large areas; it will generally not be possible to allocate all land to only one or two uses. Lower values of ϵ, for which a broader mix of land uses is required, are likely to be more applicable for real life scenarios. However, comprehensively visualising the trade-offs between objectives, as in Figure 3, can allow policy makers to better understand their options and to identify good land use plans across a number of competing goals.

7. The Glover-Martinson Method

7.1. Overview

In this section, we adapt the method of Glover and Martinson (1987) to find feasible solutions of our land use multi-objective LP that consider its multiple objective functions simultaneously. Specifically, given that the

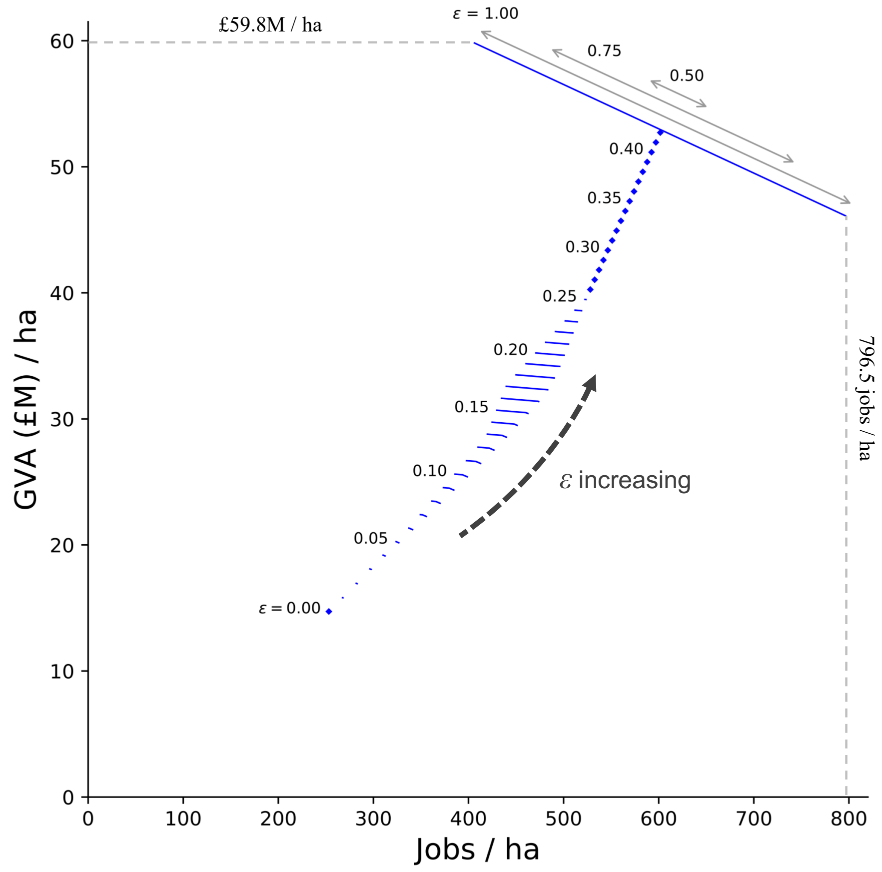

Figure 3: Pareto Frontiers of **MOLP1-P**, with respect to jobs and GVA per hectare, for increasing values of the deviation parameter, ϵ

allocation of land for housing can effectively be performed independently of the rest of our optimisation structure, we aim to find proportional land use allocations that perform well in terms of both job creation and GVA generation, using the proportional form of the model, **MOLP1-P**, discussed in Section 5.2.4.

We begin by explaining the concept and structure of the Glover-Martinson Method as applied to the general case **MOLP0**, before concentrating on the specific case of the proportional land use model **MOLP1-P**.

7.2. The Glover-Martinson Method applied to a general multi-objective LP

In essence, the Glover-Martinson Method for multi-objective optimisation attempts to find a feasible solution that is simultaneously as close as possible to the individual optima of each objective function in percentage terms. It can also accommodate weightings for the different objective functions to account for the relative priorities of a model user.

In more mathematical terms, given a general multi-objective LP of the form **MOLP0**, the method requires that the 'Single Objective Optima' $Z_{i,\max}$ are calculated and that a set of weights, $w_i \geq 0$, are defined, establishing the relative importance of the objective functions, with the default case being $w_i = 1$, $\forall i$. These weights will be considered as additional parameters, alongside the deviation parameter ϵ.

Glover and Martinson's key step is to introduce a new decision variable, which we label δ (Z in the original source), which, for a given feasible solution $\mathbf{x}$, expresses the maximum proportional divergence of any objective function from its maximum feasible value, after scaling by the chosen weights.

For example, in a scenario with two objective functions Z_1 and Z_2, each of which has an independent maximum value of 100, a feasible solution that achieved values of 97 for Z_1 and 98 for Z_2 would be associated with the value $\delta = 0.03$ (in the equal weights case), since Z_1 is 3% below its maximum, while Z_2 is only 2% below its maximum. The Glover-Martinson Method attempts to find a feasible point that minimises this maximum divergence.

This is achieved by adding the following constraints to those of the general model, **MOLP0**:

$$w_i Z_i(\mathbf{x}) + Z_{i,\max}\delta \geq w_i Z_{i,\max} \ , \quad \forall i \in \{1, \ldots, E\} \qquad (30)$$

The original objective functions are then set aside in favour of minimising the new variable δ as the sole objective function: $Z(\mathbf{x}, \delta) = \delta$.

If weights and maxima are non-zero, Inequality (30) can be rearranged to explicitly illustrate the role of δ in governing the permitted proportional divergence of each objective function from its maximum:

$$\frac{Z_i(\mathbf{x})}{Z_{i,\max}} \geq \left(1 - \frac{\delta}{w_i}\right) \tag{31}$$

We can observe from this inequality that, for example, doubling the weight of a particular objective will halve the permitted proportional divergence of that objective from its maximum feasible value, ensuring that more highly weighted objectives are more strictly constrained to take values close to their optima, as desired.

To formalise these ideas, given a particular multi-objective LP of the form **MOLP0**, the Glover-Martinson Method is to solve the following LP:

MOLP0-GM: General Glover-Martinson LP

Decision Variables:
As in **MOLP0**, along with the additional decision variable $\delta \in \mathbb{R}$, as described above.

Given/Derived Quantities and **Parameters**:
As in **MOLP0**, along with $Z_{1,\max}, \ldots, Z_{E,\max}$, the 'Single Objective Optima' of **MOLP0** (which can be considered as given quantities), and the weights $w_1, \ldots, w_E \geq 0$, as described above, considered as additional parameters.

Objective Function (to be minimised):

$$Z(\mathbf{x}, \delta) = \delta \tag{32}$$

Equality and **Inequality Constraints**:
As in **MOLP0**, along with the additional constraints defined by (30).[1]

Solution:
If it exists, the optimal value of δ is denoted as $\delta_{\min}$, with the corresponding feasible solution of **MOLP0** denoted as $\mathbf{x}_\delta$ (or the set of corresponding feasible solutions, X_δ).

If the solution is unique, it may then be used to determine the values of the objective functions of **MOLP0**:

$$Z_1(\mathbf{x}_\delta), \ldots, Z_E(\mathbf{x}_\delta) \tag{33}$$

Note that, although we have described this program in a similar format to **MOLP0** for the sake of convenience and consistency, this is in fact a *single objective* LP, since the only mathematical goal is the minimisation of δ. Therefore, unlike the multi-objective models set out earlier, this optimisation can be solved with standard LP techniques.

[1]Note that a non-negativity constraint is not required for δ, since negative values of δ would necessarily correspond to infeasible solutions (consider (30)).

*7.3. The Glover-Martinson Method applied to the proportional multi-objective land use LP, **MOLP1-P***

The previous section set out the general case of the Glover-Martinoson Method, as applied to a general multi-objective LP. We now consider the specific case where the method is applied to the Proportional Land Use Multi-Objective LP, **MOLP1-P**.

All that is required is to define appropriate weights, $w_J, w_V \geq 0$, corresponding to the objective functions $Z_{\hat{J}}$ and $Z_{\hat{V}}$, and to set out the necessary constraints of form (30):

$$w_J \hat{J}(\mathbf{q}) + \hat{J}_{\max}\delta \geq w_J \hat{J}_{\max} \tag{34}$$

$$w_V \hat{V}(\mathbf{q}) + \hat{V}_{\max}\delta \geq w_V \hat{V}_{\max} \tag{35}$$

The Glover-Martinson Method then requires that we solve the following linear program:

MOLP1-GM: Land Use Glover-Martinson LP

Decision Variables:
$q_1, ..., q_N$, as in **MOLP1-P**, along with the additional decision variable $\delta \in \mathbb{R}$, as in **MOLP0-GM**.

Given Quantities:
As in **MOLP1-P**, along with $\hat{J}_{\max}$ and $\hat{V}_{\max}$, as defined in **MOLP1-P**.

Parameters:
The deviation parameter, ϵ, as in **MOLP1-P**, along with weights $w_J, w_V \geq 0$.

Derived Quantities:
As in **MOLP1-P**.

Objective Function (to be minimised):

$$Z(\mathbf{q}, \delta) = \delta \tag{36}$$

Equality and **Inequality Constraints**:
Constraints (26)-(29), as in **MOLP1-P**, along with Constraints (34) and (35).

Solution:
$\delta_{\min}$ and $\mathbf{x}_\delta$, as defined in **MOLP0-GM** (or X_δ if no unique solution $\mathbf{x}_\delta$ exists). This solution is then used to determine the values of the objective functions of **MOLP1-P**:

$$Z_{\hat{J}}(\mathbf{x}_\delta) = \hat{J}(\mathbf{q}_\delta) \tag{37}$$

$$Z_{\hat{V}}(\mathbf{x}_\delta) = \hat{V}(\mathbf{q}_\delta) \tag{38}$$

$\mathbf{q}_\delta$ represents a possible proportional land use allocation for non-housing land, while $\hat{J}(\mathbf{q}_\delta)$ and $\hat{V}(\mathbf{q}_\delta)$ are, respectively, the corresponding number of jobs and GVA per hectare that correspond to this allocation.

8. Applying the Glover-Martinson Method to the Heathrow OA

8.1. Introduction

Having defined the Glover-Martinson approach in the context of our land use model, we are now in a position to apply the technique to the Heathrow OA, to more thoroughly explore our second research question on how to specify optimal land use allocations for job and GVA generation.

Rather than presenting a closed set of mathematical solutions to the problem, we display the results with a range of graphical visualisations, aiming to illustrate the full scope and depth of insight that can be gained through applying the Glover-Martinson method in a planning context, to optimise land use allocation.

8.2. Method

The proportional, Glover-Martinson version of the land use multi-objective LP, **MOLP1-GM**, was solved repeatedly for different values of the parameters ϵ, w_J and w_V, using the relevant values for the Heathrow OA, as set out in Section 5.2.3. Selected values of ϵ were $0, 0.01, 0.02, \ldots, 1$. Selected values of the weights were:

- $w_J = 1$, $w_V = 0$
- $w_J = 0$, $w_V = 1$
- $w_J = 1$, $w_V = 10^j$ for $j \in \{-1, -0.9, -0.8, \ldots, 1\}$

The first case is equivalent to maximising job creation while disregarding GVA; the second case is equivalent to maximising GVA generation while disregarding jobs. The third case, where both weights are non-zero, takes account of the fact that it is the relative size of the weights, rather than their absolute values, that is of relevance for the Glover-Martinson method (hence, one weight can be fixed at 1 without undermining their essential symmetry). The chosen ratios are selected uniformly on a logarithmic scale from 0.1 to 10, and were chosen after exploratory analysis to illustrate the full scope of outcomes. Note particularly that the case of equal weights $w_J = w_V = 1$ is included, when $j = 0$.

In the presentation of results, weightings will generally be specified by the quantity $\log_{10}(w_V/w_J)$, with negative values indicating a higher weighting towards job creation, and positive values indicate a higher weighting towards GVA generation. For simplicity, we take $\log_{10}(w_V/w_J) = \pm\infty$ to indicate the cases where one weight is equal to 0.

MOLP1-GM, was solved for every combination of ϵ and weighting-pair, for a total of 2323 separate optimisations.

8.3. Results and discussion

8.3.1. Maxima of GVA and jobs per hectare

As discussed in Section 6.2.3, the maximum possible jobs created per hectare is 796, which is attained when $\epsilon \geq 0.92$ and $w_V = 0$, when all land is allocated to the most efficient land use for job creation: *LU4: Business – Offices*. The maximum possible GVA per hectare is £59.8M, attained when $\epsilon \geq 0.93$ and $w_J = 0$, when all land is allocated to the most efficient land use for GVA generation: *LU2: Financial and Professional Services.*

8.3.2. GVA and job creation as ϵ, w_J and w_V vary

Figures 4 and 5 visualise the change in the optimal number of jobs and value of GVA that may be created (per hectare) under the Glover-Martinson method, for different weightings, as ϵ increases from 0 to 1. While the plots in Figure 4 show the relationship between each objective and ϵ separately, Figure 5 combines this information to visualise the trajectory of the Glover-Martinson optimum as ϵ increases, in terms of the two objectives.

The similarity between Figure 5 and Figure 3, which depicts the Pareto Frontiers of the multi-objective optimisation for different values of ϵ, is immediately obvious. Indeed the blue and green lines in Figure 5, which represent the two extreme weighting scenarios, in which one objective is assigned a weight of zero, trace out the ends of the Pareto Frontiers seen in the previous figure.

Key observations from Figures 4 and 5 are as follows:

GVA and jobs generally increase as ϵ increases.

Figures 4 and 5 show that as ϵ increases, GVA and jobs per hectare supported at the optimal land use mix also tend to increase.

This outcome is to be expected, since higher values of ϵ indicate greater freedom to diverge from the benchmark proportions, allowing the optimisation procedure to access more of the solution space. It is also consistent with what one would expect of real world land use allocation problems. The benchmark proportions reflect the current London land use mix (see Section 4.7) and, as such, they are the result of many decades of growth, both organic and planned. It is unsurprising that current land use across the city is not optimal with regard to these two specific economic objectives, given that the considerations of local governments and developers will generally be much broader and that London's land use mix was not determined by a single actor, but rather by many interacting agents over a long period.

A threshold exists at $\epsilon \approx 0.25$.

From Figure 5, we observe that the trajectories of the Glover-Martinson optima for different weightings diverge for values of ϵ from 0 to 0.25. Figure 4 demonstrates that this divergence is almost exclusively due to differences in the number of jobs created, not GVA.

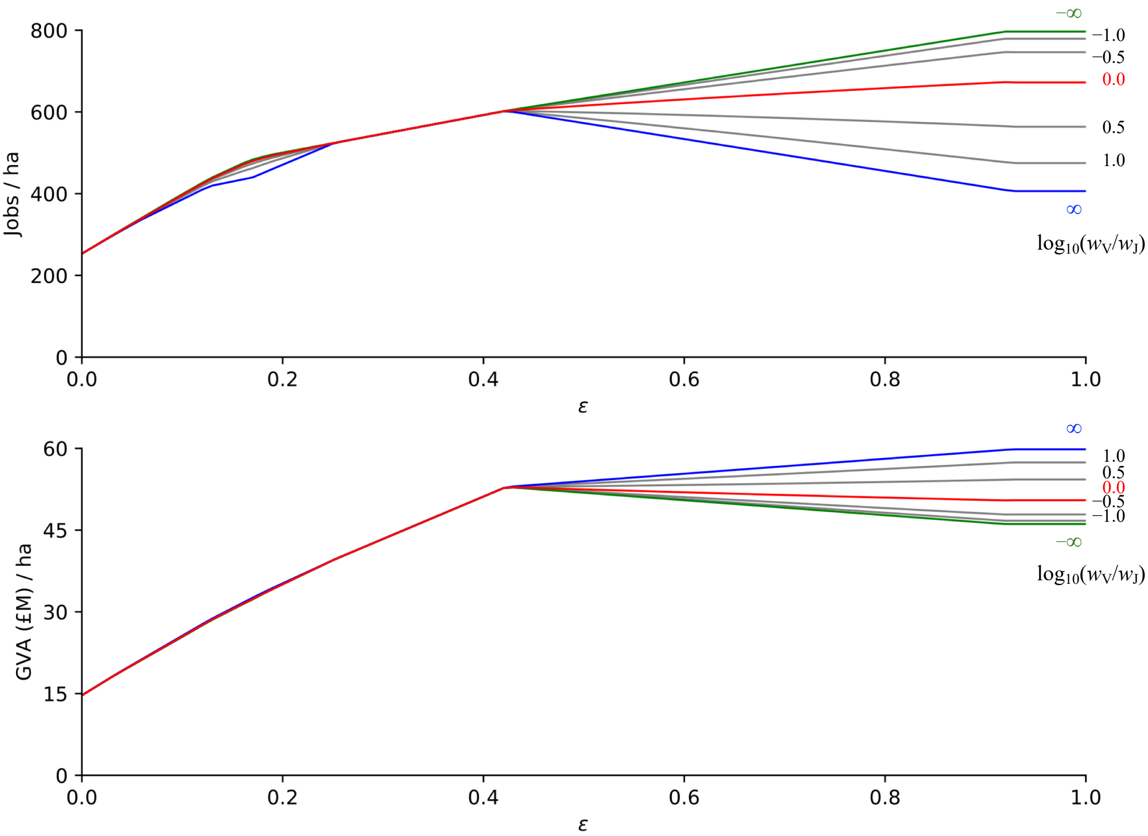

Figure 4: Line graphs representing jobs and GVA per hectare achieved at the optimal Glover-Martinson land use allocation, for various relative weightings of the two objectives, as the deviation parameter, ϵ, increases

Above the threshold at $\epsilon \approx 0.25$ (and continuing up to $\epsilon \approx 0.42$), weighting has no effect on the optimal value and the lines for all weightings converge. This convergent behaviour is easily explained, since we have already seen (in Figure 3) that, for these values of ϵ, there exists a single land use allocation that simultaneously maximises both job creation and GVA.

Relatedly, we may observe that the red line in Figure 5, which corresponds to the equal weights scenario, diverges significantly less from then green line, which corresponds to optimising exclusively on job creation, than it does from the blue line, which corresponds to optimising exclusively on GVA. This is related to the shape of the Pareto Frontiers (see Section 6.2.3), in that, for values of ϵ in this range, extremely strong GVA figures can be achieved for land use allocations lying close to the optimal allocation for job creation, but the job creation figures achieved for land use allocations close the GVA optimum are proportionally weaker. Hence, the equal weights scenario tracks the job creation optima more closely.

The divergent behaviour observed in the range $0 < \epsilon < 0.25$ may be better understood through examination of the 11 individual non-housing land uses and will be considered in more detail in Section 8.3.5.

A second threshold exists at $\epsilon \approx 0.42$**.**

For values of ϵ below 0.42, despite the divergence previously mentioned, the relative weights of the two objectives have relatively little impact on the optimal solution produced, with job and GVA generation figures differing little for different weighting schemes. However, for $\epsilon > 0.42$, this picture changes and the relative weighting of the two objectives becomes a significant determinant of the job and GVA figures that are achieved at the Glover-Martinson optimum. For high values of ϵ, therefore, trade-offs must be considered between job and GVA generation, since there there is no single land use allocation that simultaneously achieves a result close to the maxima of both objectives.

The nature of this trade-off will be examined in more detail in the context of the individual land uses in Section 8.3.5.

Clearly, these thresholds at $\epsilon \approx 0.25$ and 0.42 should be treated more as an indication of the existence of three different allocation regimes – for low, medium and high levels of restriction on land use mix – than as precise measures for real world planning scenarios. The benchmark land use mix was taken from the average land use allocation across London, so represents an indicative baseline rather than

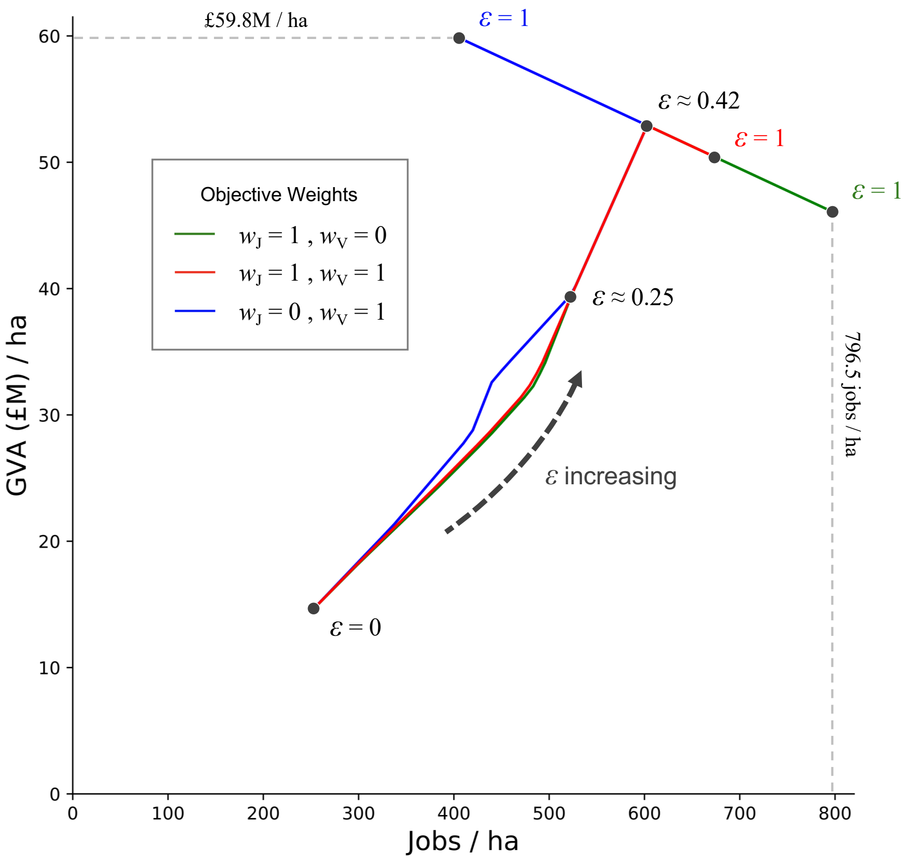

Figure 5: Trajectory of the Glover-Martinson optimum land use allocation, in terms of jobs and GVA achieved per hectare, for three different weighting schemes, as the deviation parameter, ϵ, increases

a precise one, and these thresholds and must be seen in a similar light.

When restrictions on land use mix are weak, job creation varies more widely than GVA as the weightings of the two objectives are varied.

From Figure 4, we see that for high values of ϵ, there is a significantly greater variation in jobs per hectare compared to GVA per hectare at the Glover-Martinson optima, as the relative weightings of the objectives are varied (i.e. the lines are spread more widely in Figure 4(a) than in Figure 4(b)). For values of ϵ above 0.92 (where Glover-Martinson optima are constant, for all weightings), GVA per hectare ranges from £46.1M (when $w_V = 0$) to £59.8M (when $w_J = 0$), while the jobs per hectare spread is between 406 (when $w_J = 0$) and 796 (when $w_V = 0$). In other words, if GVA is prioritised exclusively, job creation drops to around 50% of its maximum feasible value, but if job creation is prioritised exclusively, GVA only falls to just over 75% of its maximum. This suggests that, in situtations where there is significant freedom to allocate land (as represented here by high values of ϵ), prioritising job creation may be of more value than prioritising GVA, since land use allocations that support high numbers of jobs tend to also create relatively high values of GVA, while the reverse is not necessarily true.

8.3.3. Heat map matrices

The relationship between jobs per hectare, GVA per hectare, the deviation parameter ϵ, and the objective weightings w_J and w_V is also comprehensively visualised in Figure 6, which displays the Glover-Martinson optima of the objective functions for a broad range of scenarios, in the form of "heat map matrices". In each matrix, cells towards the left represent optimisations in which jobs is more heavily weighted, while cells towards the right represent optimisations in which GVA is more heavily weighted; cells towards the top represent optimisations with lower values of ϵ (more restricted land use mix) while cells towards the bottom represent optimisations with higher values of ϵ (more freedom to allocate land use).

Some of the findings from the previous section may also be observed here:

- For values of ϵ below 0.45 (approximately), the relative weighting of the two objectives has no visible impact on their optimal values. In this region, job creation and GVA increase as ϵ increases.

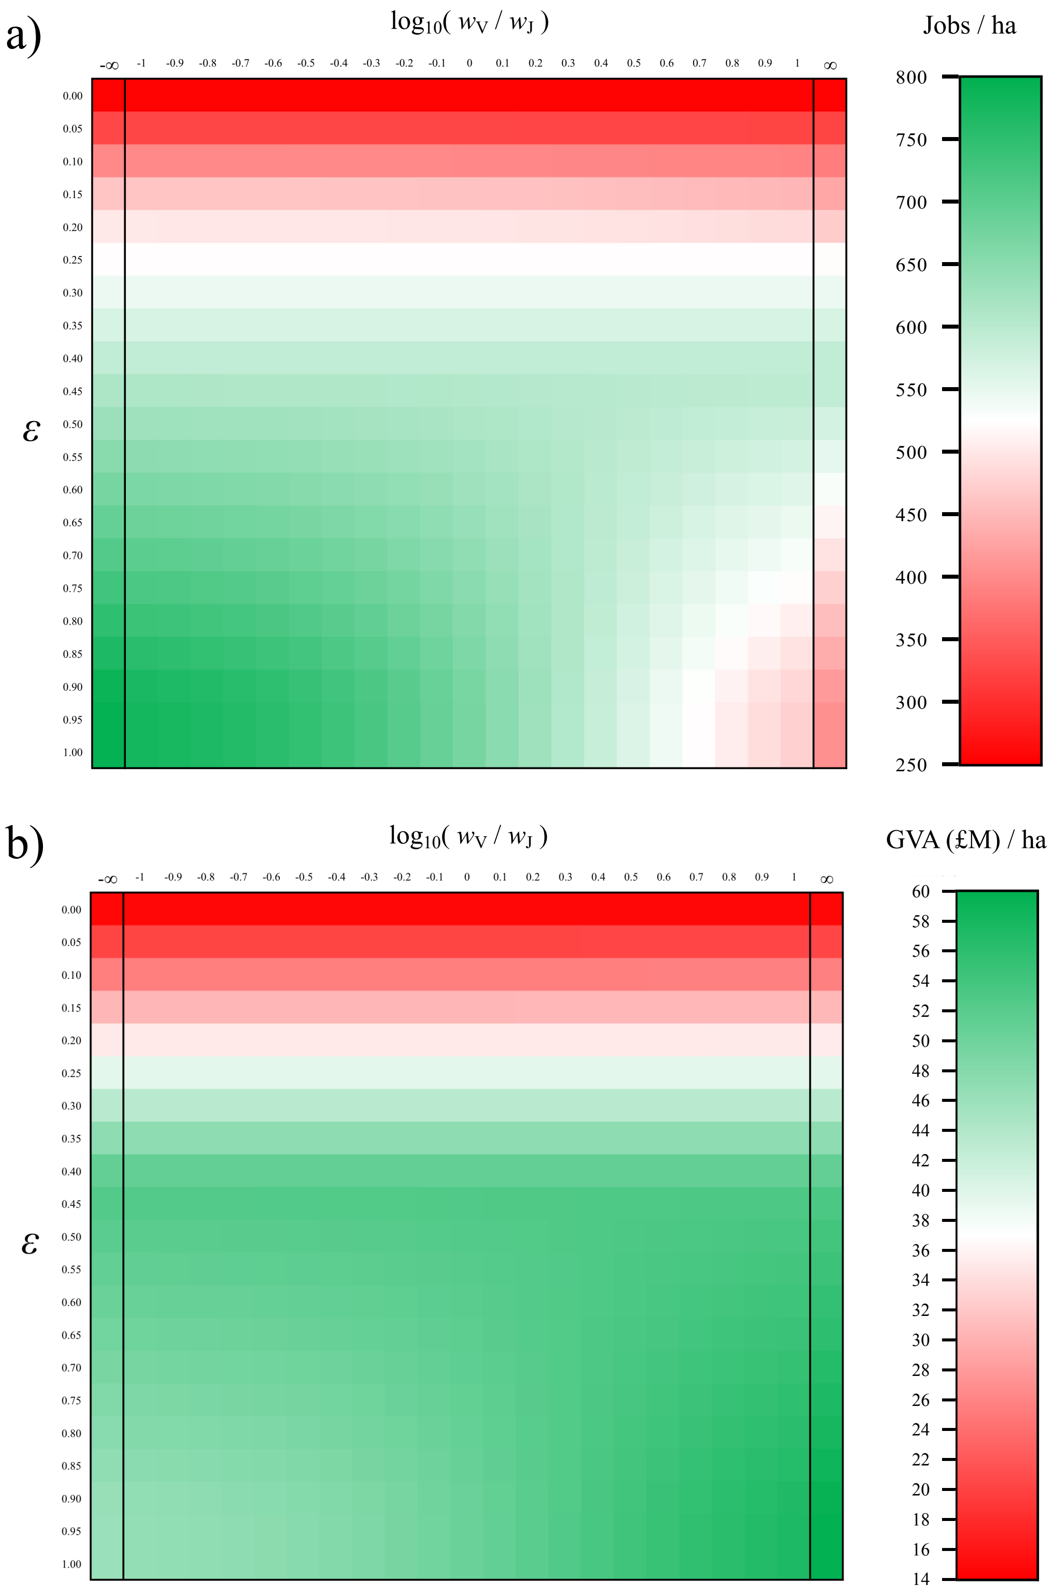

Figure 6: Heat map matrices to display jobs and GVA per hectare achieved at the optimal Glover-Martinson land use distribution, as the relative weightings of the objectives and the deviation parameter, ϵ, vary

- For higher values of ϵ, a trade-off develops between jobs and GVA, mediated by the relative weights of their respective objective functions.
- Relative to the full range of values covered by the scenarios displayed, the impact of this trade-off appears to be more severe in its effect on job creation than on GVA. The bottom right hand corner of matrix (a), representing scenarios where the GVA objective is highly weighted and ϵ is high, is seen to be associated with land use distributions that have relatively low job creation figures, but the converse effect is not observed in matrix (b).

8.3.4. Optimal land use allocations in terms of the attributes

In the previous sections, we discussed the influence of the objective weightings and the deviation parameter on the optimal values of jobs and GVA per hectare produced by the Glover-Martinson method. We now turn our attention to the characteristics of the land use allocations that give rise to these optimal values. We first discuss these allocations in terms of the three attributes discussed in Section 4.6, *Office*, *Manufacturing* and *Other Employment* (referred to as "*Other*" moving forward), since this provides an clearer setting for visualisation, communication and interpretation. We then move on to look at the 11 individual land uses in the following section.

Figure 7 presents a simplex diagram showing how land is allocated to the three attributes for $0 \leq \epsilon \leq 1$, across different weightings of the job creation and GVA objectives. The axis forming the base of the triangle shows the proportion of land allocated to *Manufacturing* land uses, the axis forming the right side corresponds to *Office* land uses, while the axis forming the left side corresponds to *Other* land uses.

The points labelled A, B and C identify features of particular importance, corresponding to the thresholds between the behavioural regimes discussed in previous sections.

Point A identifies the land use mix where all land uses are fixed at their benchmark values, corresponding to allocations of 19% of land to *Manufacturing*, 15% to *Office* and 66% to *Other*. When $\epsilon = 0$, this is the only permitted distribution of land uses.

Point B corresponds to the value $\epsilon \approx 0.25$, above which no land is allocated to the *Manufacturing* attribute, irrespective of how objectives are weighted.

Point C corresponds approximately to the range of values $0.42 \leq \epsilon \leq 1$, for which all land is allocated to the *Office* attribute, again, irrespective of how objectives are weighted.

Table 8 shows the values of ϵ and the proportion of land allocated to each of the attributes, *Office*, *Manufacturing* and *Other*, at each of these points.

Figure 8 visualises the changing land use allocations of the three attributes in a different way, demonstrating

Point	ϵ	*Office*	*Manu.*	*Other*
$\boldsymbol{A}$	0.00	0.15	0.19	0.66
$\boldsymbol{B}$	0.25	0.65	0.00	0.35
$\boldsymbol{C}$	0.42-1.00	1.00	0.00	0.00

Table 8: Specification of the points $\boldsymbol{A}$, $\boldsymbol{B}$ and $\boldsymbol{C}$, depicted in Figure 7

the relationship with ϵ more explicitly. The plots depict changes in allocation for $0 \leq \epsilon \leq 0.5$ for seven different relative weightings of the objectives, from the scenario in which GVA generation is weighted 0 at the top, to the scenario in which job creation is weighted 0 at the bottom. The x-axis is limited to $\epsilon \leq 0.5$, because after the threshold value $\epsilon \approx 0.42$, all land is allocated to the *Office* attribute.

The clearest observation from this graph is that, irrespective of the objective weightings, the proportion of land allocated to the *Office* attribute increases linearly as ϵ increases from 0 to approximately 0.42, at which point all land is allocated to this attribute. The reason for this behaviour is fairly clear, in that Table 2 clearly shows that the two land uses that make up the *Office* attribute, *LU2: Financial and Professional Services* and *LU4: Business – Offices*, produce more jobs and more GVA per hectare than any other land use.

This observation suggests that, in terms of job creation and GVA, irrespective of how these two objectives are weighted against each other, as much non-housing land as possible should be allocated to *Office* land uses.

Returning to Figure 8, we see that the way that land is distributed between the other two attributes, *Manufacturing* and *Other*, does differ slightly according to the chosen objective weights. In all cases, the share of land allocated to the *Manufacturing* attribute ultimately declines to zero as ϵ increases, but when GVA is assigned relatively higher weightings, this decline is preceeded by a regime (from $\epsilon = 0$ up to around $\epsilon = 0.2$) in which the share allocated to *Manufacturing* is relatively stable. Conversely, when job creation is assigned higher weightings, land allocated to the *Other* attribute is initially fairly stable (from $\epsilon = 0$ up to around $\epsilon = 0.1$) before declining, but this decline becomes more immediate as the weight of GVA is increased.

This observation suggests that *Manufacturing* land uses make up a more significant component of land use allocation in situations where land use mix is fairly restricted and where GVA is prioritised over job creation.

8.3.5. Allocation to individual land uses

Figures 9 and 10 are similar to Figure 8, considering the same seven relative weightings of the objective functions, but they portray changes in all 11 individual land uses rather than the three attributes. Figure 9 covers the full range of values $0 \leq \epsilon \leq 1$, clearly depicting the trade-off between the two land uses that make up the *Office* attribute, which occurs above the threshold $\epsilon \approx 0.42$, while Figure 10 zooms in to display the range $0 \leq \epsilon \leq 0.25$, in

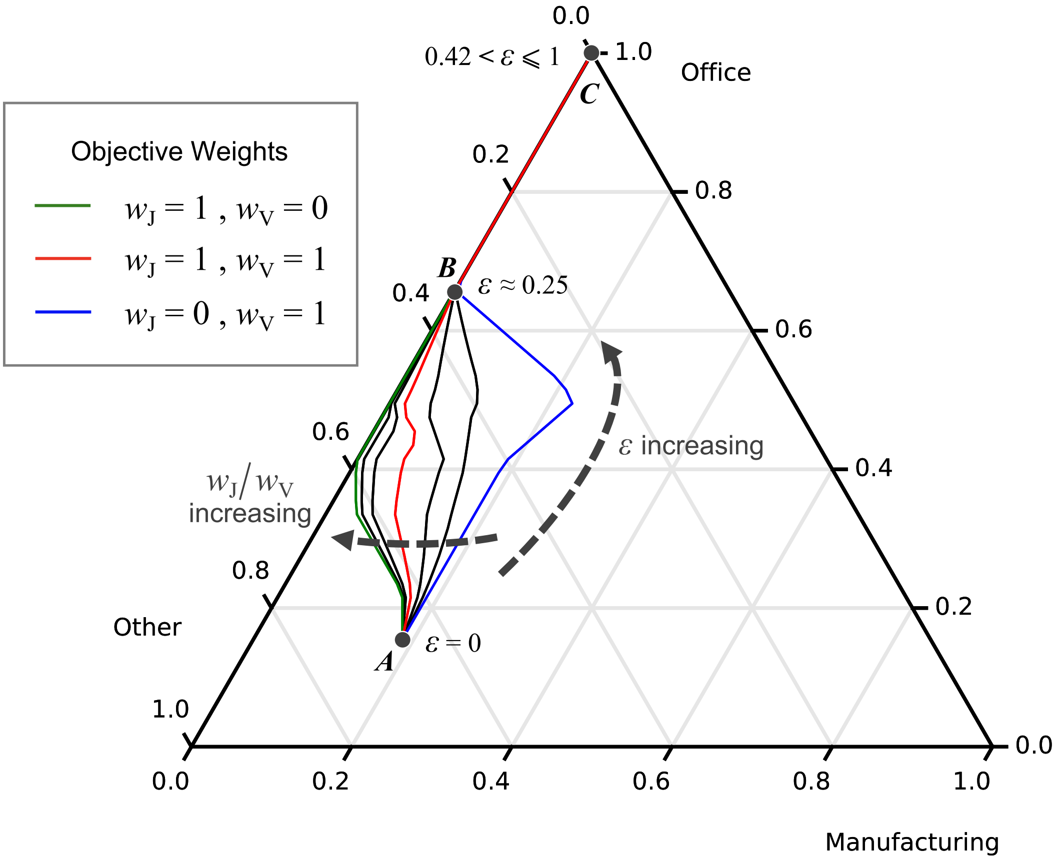

Figure 7: Simplex plot depicting the trajectory of the optimal Glover-Martinson land use allocation, in terms of the three land use attributes defined in Table 3, for various relative weightings of the two objectives, as the deviation parameter, ϵ, increases

which complex exchanges between other land uses can be observed.

The three regimes of behaviour identified in previous sections can also be observed in Figure 9. We begin by considering three particular land uses – *LU2: Financial and Professional Services* and *LU4: Business – Offices* (the two *Office* land uses) and *LU1: Shops* – and how their allocations change in each regime.

The first regime is $0 \leq \epsilon \leq 0.25$ (between Point A and Point B in Figure 7), and is visualised more clearly in Figure 10. For values of ϵ in this range, for all weighting combinations, many land uses are active in the Glover-Martinson optima (i.e. assigned to a non-zero proportion of land), since there are still significant restrictions on how land may be distributed relative to the benchmarks. One clear observation from Figure 9 is that our three selected land uses increase their proportions linearly over this range, for all weighting combinations.

We had already observed that *Office* land uses should be prioritised, but this observation suggests that retail land will also make up an important part of the land use distribution in situations where land use mix is somewhat restricted.

The second regime is where ϵ lies between 0.25 and 0.42 (between Point B and Point C in Figure 7). For these values of ϵ, only the three land uses mentioned above are active.

Figure 9 confirms that the optimal land use allocation proposed by the Glover-Martinson Method in this regime is wholly independent of the objective weightings. Indeed, the only change in land use allocation is that the proportion of land allocated to *LU1: Shops* declines linearly to zero, while the two *Office* land uses continue to linearly increase, as before.

The third and final regime is where $\epsilon > 0.42$ (Point C in Figure 7). As we have seen, by this point, all land is allocated to the two *Office* land uses. The only change is how land is allocated between them, with similar proportions allocated to each at $\epsilon = 0.42$ and growing divergence as ϵ increases.

The nature of this divergence depends strongly on the the weightings of the GVA and job creation objectives. When job creation is preferred, more land is allocated to *LU4: Business – Offices*. Conversely, when GVA generation is weighted higher, more land is allocated to *LU2: Financial and Professional Services*. In the case of equal weights, *LU4: Business – Offices* is favoured somewhat.

This observation suggests that, where there is very sig-

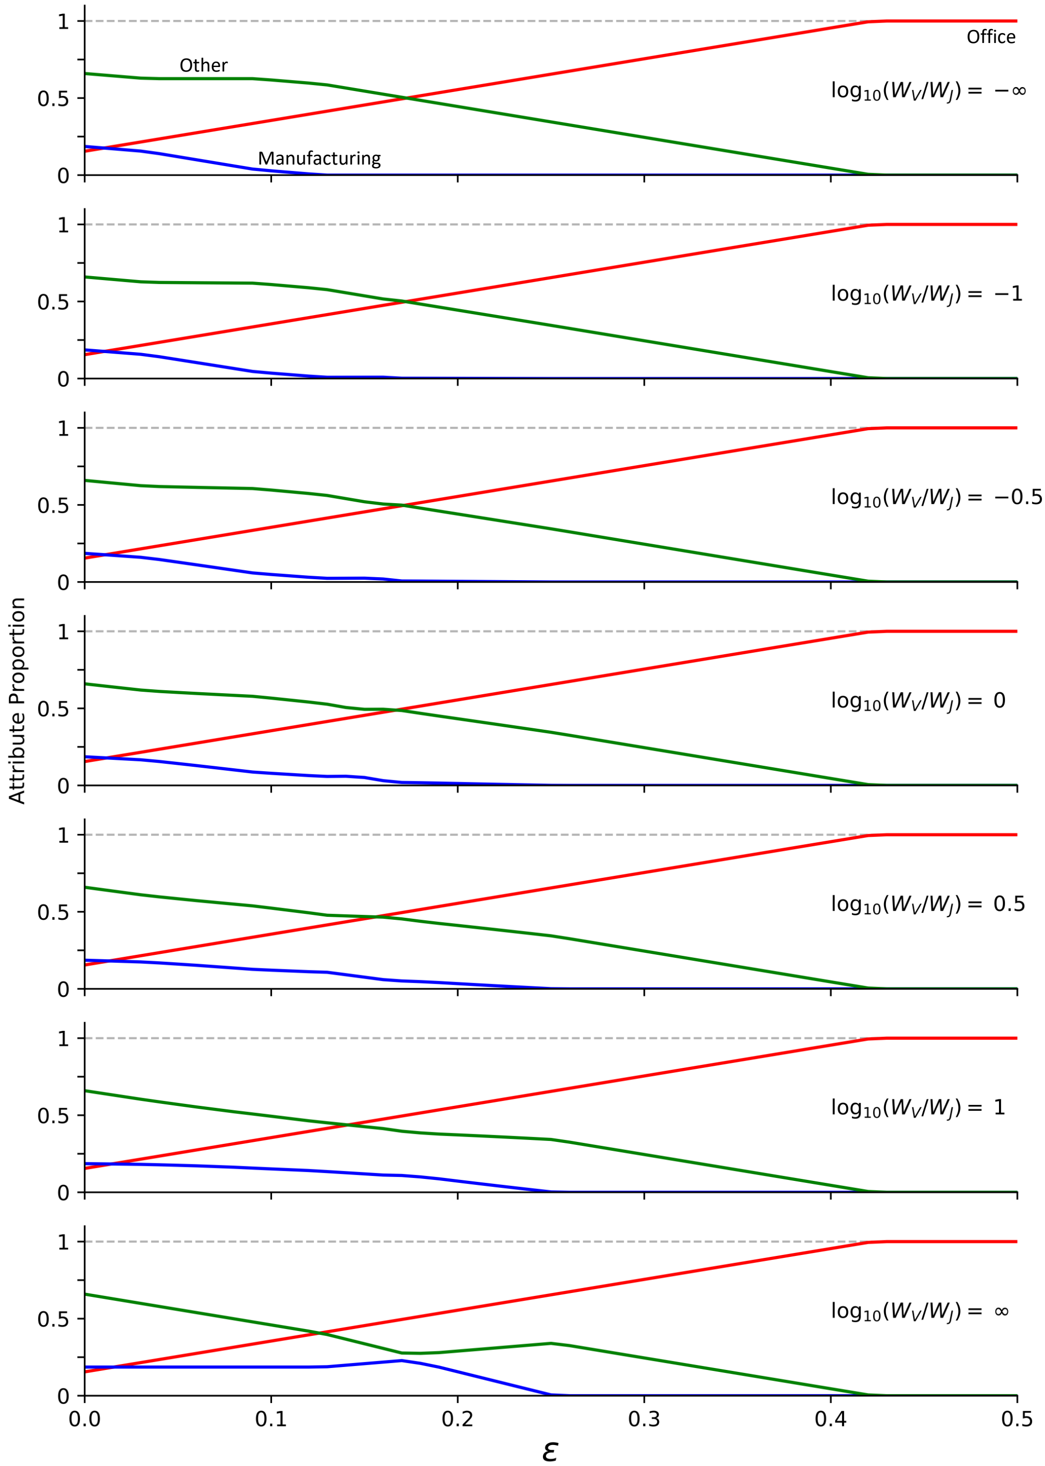

Figure 8: Line graphs depicting the proportion of land allocated to each of the three attributes at the optimal solutions produced by the Glover-Martinson Method, for varying weightings of job and GVA creation, for $0 < \epsilon < 0.5$

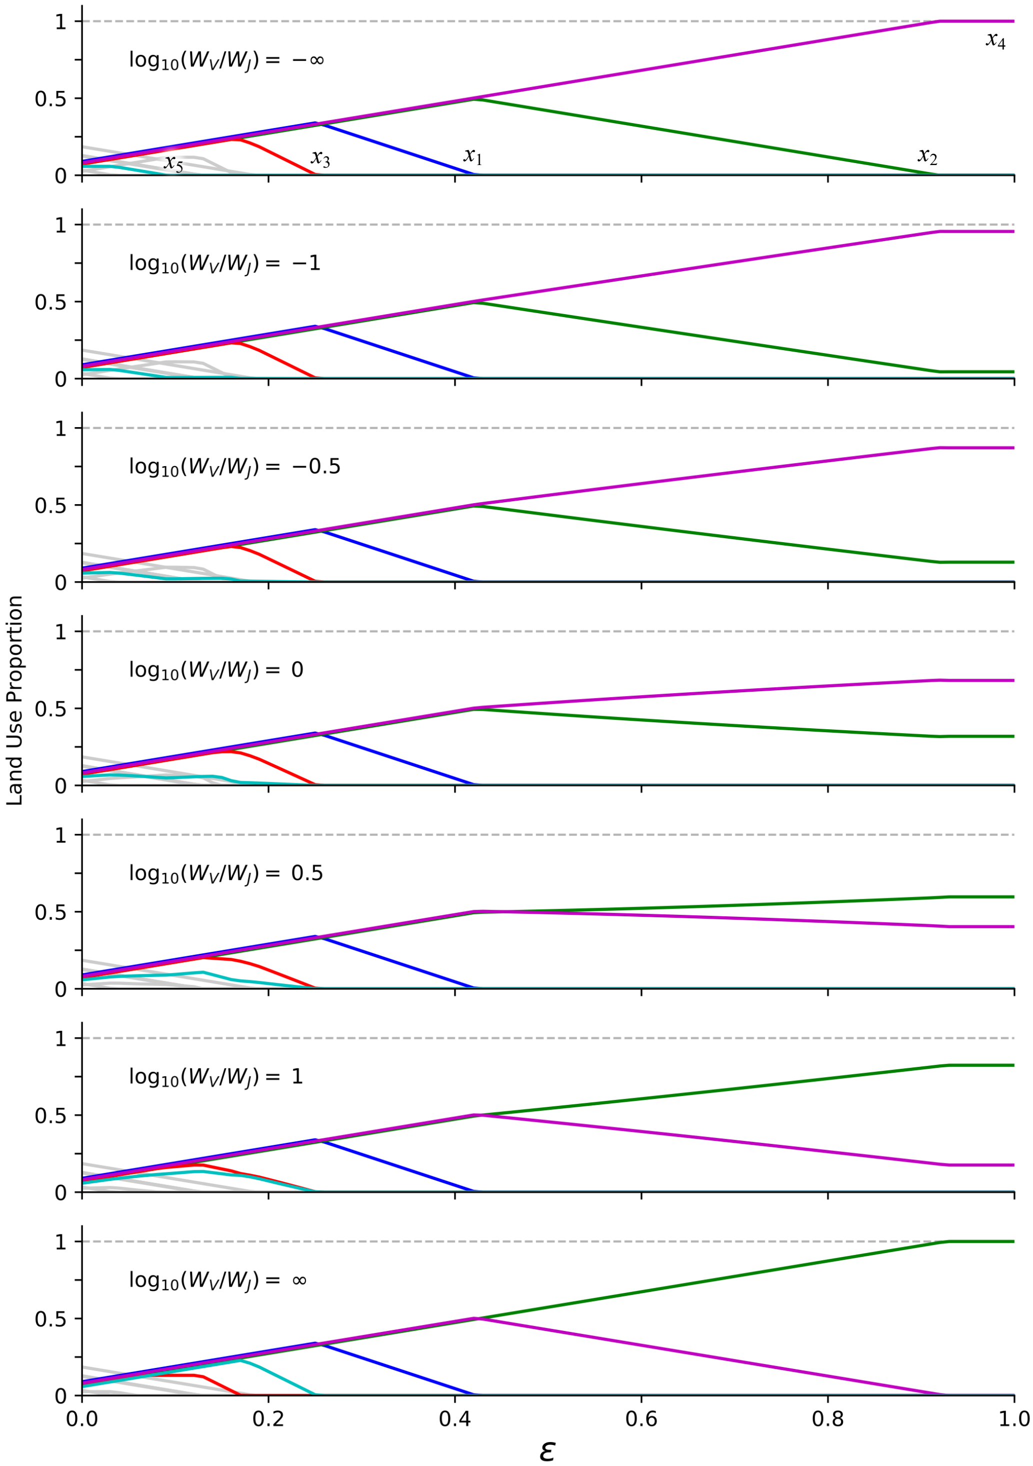

Figure 9: Line graphs depicting the proportion of land allocated to each of the eleven selected land uses at the optimal solutions produced by the Glover-Martinson Method, for varying weightings of job and GVA creation, for $0 < \epsilon < 1$. Note that land uses 6-11 are all represented in pale grey.

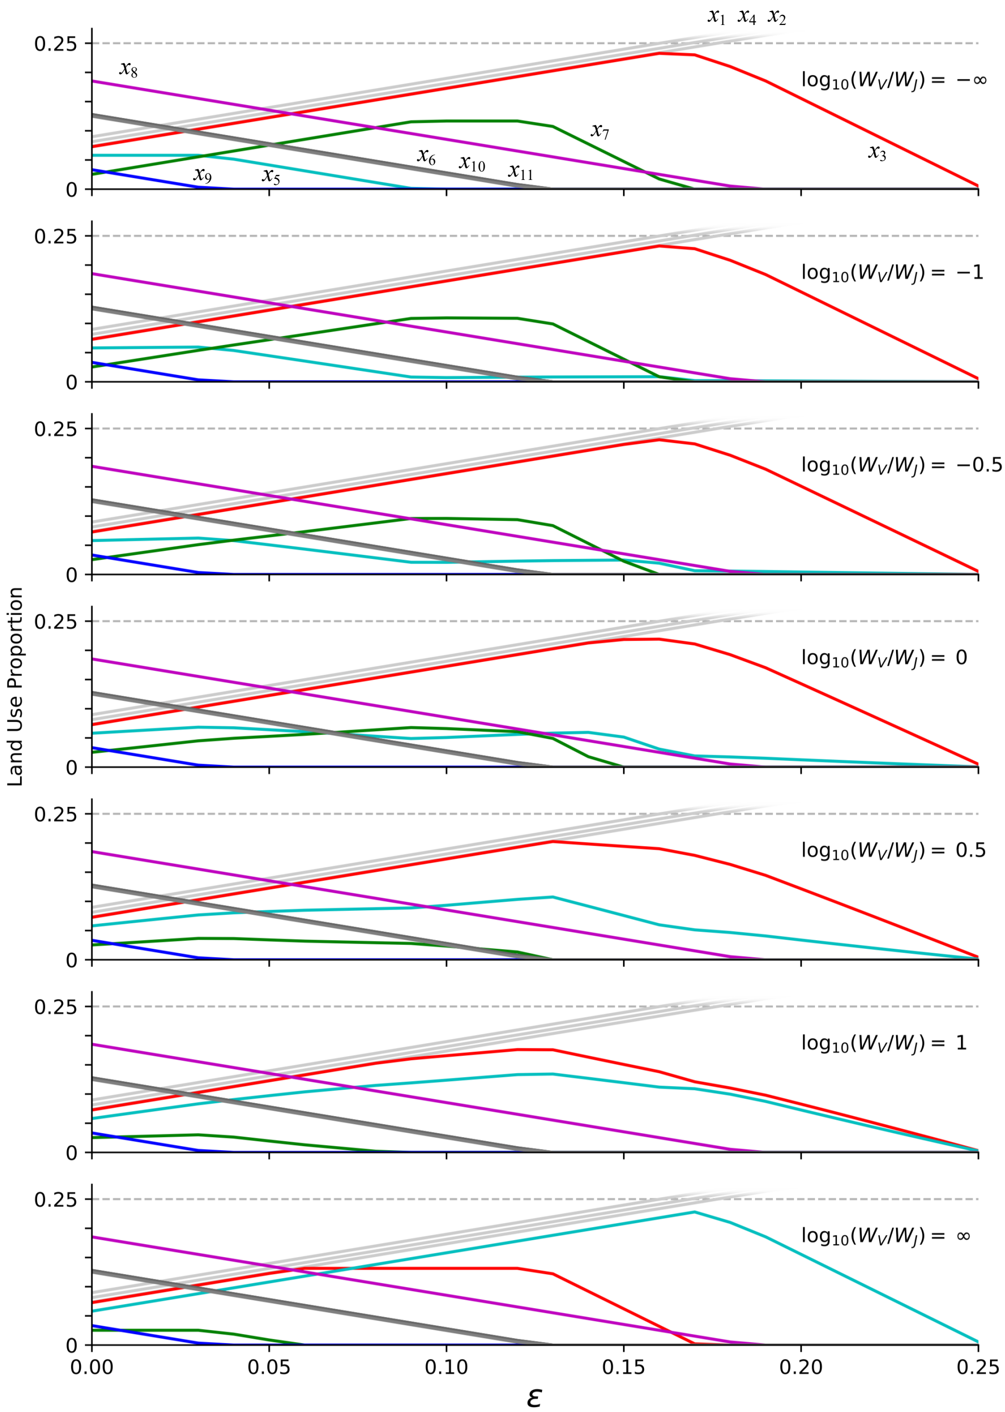

Figure 10: Line graphs depicting the proportion of land allocated to each of the eleven selected land uses at the optimal solutions produced by the Glover-Martinson Method, for varying weightings of job and GVA creation, for $0 < \epsilon < 0.25$. Note that land uses 1, 2 and 4 are all represented in pale grey, while the lines representing land uses 6, 10 and 11, represented in darker greys, are almost collinear in all plots.

nificant freedom to allocate land, whether to favour *LU4: Business – Offices* or *LU2: Financial and Professional Services* depends on the relative importance assigned to the objectives of job creation and GVA. If prioritising GVA, the latter land use is a better option; if prioritising job creation, the former is preferable. More balanced priorities would be better addressed with a mixture of the two land uses, with somewhat more emphasis on the former.

The divergence between these two land uses clarifies the trade-off between job creation and GVA that was identified in previous figures. It explains the widening Pareto Frontiers in the top right corner of Figure 3, the divergence of trajectories in Figures 4 and 5, and so on.

It remains to consider the remaining eight land uses and how their allocation responds to changes in ϵ and the objective weightings, as visualised in Figure 10. We note that, for values of ϵ above the threshold of 0.25, none of these land uses is active in any scenario.

From the figure, we see that the proportions allocated to five land uses drop linearly to 0 as ϵ increases, irrespective of the weightings assigned to job and GVA creation: *LU6: Storage or Distribution*, *LU8: Residential Institutions*, *LU9: Non-Residential Institutions*, *LU10: Assembly and Leisure*, and *LU11: Sui Generis*. This suggests that a minimal amount of land should be allocated to these uses if attempting to optimise job and GVA creation.

Of the remaining three land uses, the most prominent in Figure 10 is *LU3: Restaurants, Cafes, Bars, Takeaways*. The allocation of land to this use initially increases linearly with ϵ, peaking between $\epsilon \approx 0.13$ and $\epsilon \approx 0.17$ (depending on the weightings), then dropping to 0. Where job creation is equally or more highly weighted than GVA, over 20% of land is allocated to this single use at its peak, though its contribution is somewhat lower when GVA generation is more heavily weighted.

LU7: Hotels has a similar pattern, though it makes up a smaller part of the land use mix for all weighting schemes. It is allocated over 10% of available land at its peak, when ϵ is around 0.1 and job creation is more heavily weighted; otherwise, this land use plays a fairly minor role in Glover-Martinson optima.

The *LU5: General Industrial* land use has the opposite pattern from the two above, forming a significant part of the land use mix only when GVA is heavily weighted. Its peak of over 20% is attained (at $\epsilon \approx 0.17$) only when job creation is allocated a weighting of zero. For the equal weighting case, it is allocated a fairly stable proportion of just over 5% of available land for values of ϵ below 0.15, before dropping off, and it makes up a diminishing proportion of allocated land when job creation is weighted more highly.

These observations suggest that *LU7: Hotels* and particularly *LU3: Restaurants, Cafes, Bars, Takeaways* can form a very significant part of the land use mix when jobs are prioritised. When GVA is strongly prioritised, however, the *LU5: General Industrial* land use is a more favourable option.

8.4. Summary and implications

8.4.1. Summary of findings from the Glover-Martinson approach

The Glover-Martinson Method has strengthened and deepened some of the conclusions drawn from an inspection of the Pareto Frontiers, depicted in Figure 3.

A summary of the findings is as follows.

The optimisations conducted indicated three distinct regimes of land use allocation, depending on the values of our parameters.

The first regime relates to situations in which land use mix is strictly to moderately restricted, as represented by values of ϵ between 0 and approximately 0.25. In this regime, optimal land use allocations specified by the Glover-Martinson Method involve a wide range land uses. Based on the attributes defined in Section 4.6, optimal solutions in this regime therefore tend to feature a mixture of all three: *Manufacturing*, *Office* and *Other* land uses.

In this regime, changing the relative weightings of job creation and GVA generation does have an impact on the optimal solutions produced for certain land uses, particularly *LU3*, *LU5* and *LU7*. However, the impact of these differences on the number of jobs produced and, more especially, the GVA generated are small. Optimal job creation figures per hectare in this regime increase from approximately 250 (at the benchmark land use proportions, $\epsilon = 0$) to 520 (at $\epsilon \approx 0.25$). Optimal GVA generation figures per hectare increase from approximately £15 M to £39 M.

The second regime relates to situations in which land use mix is much less strictly restricted, as represented by values of ϵ between 0.25 and approximately 0.42. In this regime, no land is allocated to the *Manufacturing* attribute. Indeed, all land is allocated to just three land uses: the two *Office* land uses, *LU2* and *LU4* along with *LU1: Shops*. Allocation is very simple, in that land is assigned roughly equally to the three land uses when $\epsilon \approx 0.25$, with *LU1* declining steadily to zero as ϵ increases, while *LU2* and *LU4* increase steadily to an approximately even split by $\epsilon \approx 0.42$.

In this regime, changing the relative weightings of job creation and GVA generation has no impact on the optimal solutions produced; a single land use allocation optimises both job creation and GVA generation. Optimal job creation figures per hectare in this regime increase from approximately 520 (at $\epsilon = 0.25$) to 600 (at $\epsilon \approx 0.42$). Optimal GVA generation figures per hectare increase from approximately £39 M to £53 M.

The third regime relates to situations in which land use mix is only very loosely restricted, as represented by values of ϵ between approximately 0.42 and 1. In this regime, all land is allocated to the the two *Office* land uses, *LU2* and *LU4*.

This is the only regime in which the relative weightings of job creation and GVA generation have a significant impact on both the optimal land use allocations produced and the actual job and GVA creation figures realised. There is a direct trade-off between *LU4: Business – Offices*, the best land use for job creation, and *LU2: Financial and Professional Services*, the best land use for GVA generation.

Maximising only for job creation, for the highest values of ϵ, results in all land being allocated to *LU4*, with corresponding job and GVA figures per hectare of approximately 796 and £46 M. Maximising only for GVA generation, for the highest values of ϵ, results in all land being allocated to *LU2*, with corresponding job and GVA figures per hectare of approximately 406 and £60 M. The case of equal weights, for the highest values of ϵ, produces an allocation favouring *LU4* (68%) over *LU2* (32%), corresponding to job and GVA figures per hectare of approximately 672 and £50 M.

8.4.2. Implications for the Heathrow OA

A brief interpretation of the above findings, along with earlier findings relating to housing provision, in terms of the Heathrow OA, is as follows.

Regarding housing, a desired density, lying in the range 35 to 170 units per hectare, should be established and sufficient land should be earmarked from the available 700 ha to meet the target given the chosen density. A higher target could indeed be set, since there is plenty of leeway to meet the existing targets for jobs and homes in an area of this size.

It is likely that a broad mixture of land uses will be required from the remaining hectares, meaning that our scenarios with lower values of ϵ are of most relevance for the allocation of non-housing land. Within the restrictions on land use mix, however, the following considerations should be made.

Relative preferences between job and GVA creation are of little relevance, since good land use allocations should produce strong results on both of these measures if land use mix is fairly restricted. Given this observation, as much land as possible should be allocated to *LU1: Shops*, *LU4: Business – Offices* and *LU2: Financial and Professional Services*. Simultaneously, as little land as possible should be allocated to *LU6: Storage or Distribution*, *LU8: Residential Institutions*, *LU9: Non-Residential Institutions*, *LU10: Assembly and Leisure*, and *LU11: Sui Generis*. The remaining three land uses, *LU3: Restaurants, Cafes, Bars, Takeaways*, *LU5: General Industrial* and *LU7: Hotels* can be allocated more or less land, as desired.

Obviously, these conclusions relate only to the two economic objectives under consideration (and to satisfying the housing target). The true planning context of the Heathrow OA involves many more factors and restrictions than are covered by this framework. However, within the bounds of any such additional conditions and goals, these findings provide an indication to planners of how job and GVA creation can be favoured.

Section 9 will discuss further limitations and provisos for these findings, and place them in a broader context.

9. Conclusions and further work

Our investigation has produced the following answers to our research questions.

1. Are the latest guideline figures for homes and jobs in the Heathrow OA achievable and to what extent could development be concentrated on brownfield sites, avoiding development on Green Belt land?

The required figures for home and job creation, as given in the most recent versions of the London Plan, are comfortably achievable in the Heathrow OA, based on the stated area of 700 ha. However, the amount of available brownfield land in the OA is insufficient to accommodate the necessary development. Green Belt land will almost certainly need to be included in any development plan if the figures are to be achieved.

2. How might optimal allocations of land use in the Heathrow OA be specified in order to maximise home and job creation, while also maximising the gross value added (GVA) supported by the available land?

Consideration of Pareto Frontiers and application of the Glover-Martinson Method for multi-objective optimisation has provided a comprehensive picture of how land use might be allocated to achieve these goals. These results are summarised in Section 8.4.

A secondary goal of this piece has been to illustrate the value of multi-objective linear programming as a tool for planners, highlighting the many different perspectives that the method can provide and presenting these perspectives in visual and intuitive ways. Despite limiting ourselves to two basic approaches from among the many discussed in the literature review (Pareto Frontiers and the Glover-Martinson Method), a deep and rounded picture of the economic trade-offs of our planning scenario has been possible.

However, we would emphasise that linear programming approaches should ultimately be considered simply as a tool to support planning decisions. We make no claim that they provide a single "perfect" solution for any land use allocation problem. Indeed, we have observed that real life planning problems will invariably involve factors that cannot be easily translated into a mathematical framework. It is for this reason that we have used optimisation techniques to provide a suite of possible land use allocations and to devise the broad characteristics of a good land use allocation, rather than offering a single solution. Maintaining this flexibility should allow planners to simultaneously quantify those aspects of a problem that can be quantified, while also applying their own judgement and experience to select a particular solution.

Real life planning problems will generally involve many more restrictions than those that have been applied in our model, not least on the degree of control that planners actually have over how land is ultimately used. While it may be possible for a planning authority to incentivise the development of particular forms of industry within an area, we are fully aware that it will generally not be possible to completely specify the exact land use of every hectare of land within a development to the extent that we have categorised land use here. This is a further reason to maintain flexibility in the outputs of land use optimisation models.

One particular area that requires further consideration is how to guarantee realistic land use mix in the distributions proposed by land use models. In our model, we used the devices of benchmark proportions and the deviation parameter to ensure that model outputs featured a diverse range of land uses, rather than allocating all land to one or two of the most efficient categorisations. However, this approach is clearly a very significant simplification of the reality of how land use is allocated in real planning cases. The ideal situation would be that a realistic land use mix arises naturally from the detailed constraints of a particular problem. However, we consider that a more general treatment of land use mix in multi-objective LP is an open question, necessitating further research.

Some further notes of caution should be raised about how to interpret our results. There is a difference between calculating the jobs and GVA generated by existing land assigned to particular land uses and assuming that allocating new land to those same land uses will produce the same amounts of jobs and GVA. There are clear considerations of supply and demand to bear in mind. This is a further reason to underline that results presented here should be seen as indicative and not precise. There is a significant amount of unquantifiable uncertainty over the job and GVA figures that will be produced by any particular development plan, and outcomes will also be highly dependent on future economic conditions. That being said, we believe that clear quantitative approaches like the multi-objective LP methods presented here are of value in providing a "most-informed estimate" to planners, to support their decision making.

It should also be noted that our findings are dependent on three key figures provided in the London Plan for the Heathrow OA: the guideline figures for homes and jobs, and the stated amount of available land. Despite the apparent lack of progress with regard to adopting a planning framework for the Heathrow OA, each subsequent version of the London Plan has maintained or increased the guideline figures for homes and jobs.

Since no detailed explanation has been provided for how any of these figures were derived, it is reasonable to question whether they are reasonable and, indeed, what precisely they mean. For example, it does not seem credible that any development plan for the area could involve construction on a full 700 ha of land, so the effective amount of land available may be significantly lower than this stated figure, implying that there could be considerably less leeway to achieve the targets than our analysis suggests.

Furthermore, since the publication of the first London Plan in 2004, over 23,500 new homes have been built in the boroughs of Hillingdon and Hounslow (MHCLG, 2019). Although the Heathrow OA does not fully encompass these two boroughs, it is likely that a certain number of homes have been constructed within the OA boundaries. However, it is unclear whether these housing completions have been taken into account in setting the most recent targets. Also, given that additional development has occurred in the interim, we would question whether some of the land that had initially been earmarked for homes and jobs may no longer be available.

Further work is planned on both the application and the theory of multi-objective LP for land use optimisation. This will include technical adaptations of the Glover-Martinson Method and applications of the method of Makowski et al. (2000) to the Heathrow OA. It would also be valuable to integrate some of the more explicitly spatial approaches to land use LP that were discussed in the Literature Review, allowing for a spatial plan of where development in the Heathrow OA could actually be concentrated, accounting for transportation links, the relative locations and interactions between different land uses, and so on. However, the necessary data to support such an approach does not appear to be publicly available, so work of this nature would necessitate closer work with planning authorities or a significant data collection effort.

Spatially explicit land use data would allow for the construction of models that better account for synergies between new and existing land uses. For example, it may make sense to allocate more land in the OA to transport and logistics land uses given the importance of Heathrow Airport, even if such an allocation would not be considered optimal given our objectives. Similarly, it could be reasonable to question whether there any additional housing should be constructed in the immediate vicinity of the airport, given the evident issues of noise and air pollution.

While the motivation behind the designation of OAs was to localise the broader London Plan and to plot a more specific path to achieving its wider goals, it is not necessarily the case that seeking optimal land use plans at the OA scale represents the most effective use of our methods. Conducting a multi-objective LP analysis at a city-wide scale could also be valuable for the consideration of local synergies and economies of scale. Comparing locally optimal and city-wide optimal land use allocations could be another interesting avenue of research.

As a final point, with respect to the guideline figures for homes and job creation set out in all editions of the London Plan, we would suggest that they may lack the nuance required to effectively address London's housing and employment needs. A serious issue with housing provision in London is the lack of affordable housing, so the possibility of defining specific targets for affordable homes should be considered. Similarly, job creation targets could be disaggregated by sector or wage bracket, to specify the types of jobs that are desirable or necessary in a certain area. This approach could be used to address skills shortages and link with a broader industrial strategy. Setting such detailed targets could have a genuine impact on the utility of future development, allowing for a more informed approach to future work on land use optimisation and, consequently, leading to more effective solutions to London's issues.

Bibliography

Aerts, J.C.J.H., E. E. H. G., Stewart, T., 2003. Using linear integer programming for multi-site land-use allocation. Geographical Analysis 35 (2), 148–169.

Alexander, C., 1964. Notes on the Synthesis of Form. Harvard University Press, Cambridge MA.

Arthur, J. L., Nalle, D. J., 1997. Clarification on the use of linear programming and GIS for land-use modelling. International Journal of Geographical Information Science 11 (4), 397–402.

Aviation Environment Federation, 2019. CCC Net Zero Report: We'll still be flying in 2050, but Government can no longer ignore aviation emissions in its climate policies [press release].

Bammi, D., Bammi, D., Feb 1979. Development of a comprehensive land use plan by means of a multiple objective mathematical programming model. Interfaces 9 (2), 50–63.

BBC, 2010. Heathrow runway plans scrapped by new government. *BBC News* [online] 12 May 2010. Available at: news.bbc.co.uk [Accessed 25 Nov 2019].

BBC, 2018. Heathrow airport: MPs vote in favour of expansion. *BBC News* [online] 26 Jun 2018. Available at: news.bbc.co.uk [Accessed 25 Nov 2019].

CAG Consultants, 2016. London Employment Sites Database: Final Report. London: CAG Consultants.

Chuvieco, E., 1993. Integration of linear programming and gis for land-use modelling. International Journal of Geographical Information Science 7 (1), 71–83.

Committee on Climate Change, May 2019. Net Zero – Technical Report. London: Committee on Climate Change.

Day, J., 1973. A linear programming approach to floodplain land use planning in urban areas. American Journal of Agricultural Economics 55 (2), 165–174.

Delgado-Matas, C., Pukkala, T., 2014. Optimisation of the traditional land-use system in the angolan highlands using linear programming. International Journal of Sustainable Development & World Ecology 21 (2), 138–148.

Department for Transport, June 2005. Annual Report (CM 6527). London: UK Government.

Department for Transport, 2018. Airports National Policy Statement: New runway capacity and infrastructure at airports in the South East of England.
URL `https://assets.publishing.service.gov.uk`

Department for Transport, Grayling, C., 2016. Government decides on new runway at Heathrow.
URL `https://www.gov.uk/`

Downey Brill, Jr., E., May 1979. The use of optimization models in public-sector planning. Management Science 25 (5), 413–422.

Eurostat, 2019. Glossary: Gross value added.
URL `https://ec.europa.eu/eurostat/`

GLA, 2004. The London Plan.
URL `https://www.london.gov.uk`

GLA, 2006. Housing space standards.
URL `https://www.london.gov.uk`

GLA, 2008. The London Plan 2008 (consolidated with alterations since 2004).
URL `https://www.london.gov.uk`

GLA, 2011. The London Plan 2011.
URL `https://www.london.gov.uk`

GLA, 2015a. GVA per workforce job in London and the UK. [online] Available through: London Datastore https://data.london.gov.uk.

GLA, 2015b. Ward profiles and atlas.
URL `https://data.london.gov.uk/`

GLA, 2016a. Economic evidence base for London 2016.
URL `https://www.london.gov.uk`

GLA, 2016b. The London Plan.
URL `https://www.london.gov.uk`

GLA, 2017. The London Plan - draft for public consultation.
URL `https://www.london.gov.uk`

GLA, 2018. London Plan opportunity areas.
URL `https://data.london.gov.uk/`

GLA, 2019a. London Plan annual monitoring report.
URL `https://www.london.gov.uk`

GLA, 2019b. London plan opportunity areas.
URL `https://data.london.gov.uk/`

GLA, 2019c. The draft London Plan.
URL `https://www.london.gov.uk/`

Glover, F., Martinson, F., Mar 1987. Multiple-use land planning and conflict-resolution by multiple objective linear-programming. European Journal of Operational Research 28 (3), 343–350.

Hayton, K., Feb 1981. A linear programming land selection model for structure planning: A case study of tyne and wear. Regional Studies 15 (6), 425–437.

Heathrow, Jan 2018. Heathrow expansion consultation: Our emerging plans.

Heathrow, 2019. Facts and figures.
URL `https://www.heathrowexpansion.com`

Henseler, M., Wirsig, A., Herrmann, S., Krimly, T., Dabbert, S., Apr 2009. Modeling the impact of global change on regional agricultural land use through an activity-based non-linear programming approach. Agricultural Systems 100 (1-3), 31–42.

Kumar, P., Rosenberger, J. M., Iqbal, G. M. D., Dec 2016. Mixed integer linear programming approaches for land use planning that limit urban sprawl. Computers & Industrial Engineering 102, 33–43.

Liu, Y., Qin, X., Guo, H., Zhou, F., Wang, J., Lv, X., Mao, G., Dec 2007. ICCLP: An inexact, chance-constrained linear programming model for land use management of lake areas in urban fringes. Environmental Management 40 (6), 966–980.

LPA, 2018. Brownfield land register.
URL `https://data.london.gov.uk/`

LWT, June 2018. Wildlife sites threatened by government's 'yes' to Heathrow Airport expansion.
URL `https://www.wildlondon.org.uk`

MacGregor Smith, J., Liebman, J. S., Jun 1978. A zero-one integer-programming formulation of the problem of land-use assignment and transportation-network design. Environment and Planning B 5 (1), 101–115.

Makowski, D., Hendrix, E., van Ittersum, M., Rossing, W., Jun 2000. A framework to study nearly optimal solutions of linear programming models developed for agricultural land use exploration. Ecological Modelling 131 (1), 65–77.

MHCLG, 2018. English local authority green belt dataset.
URL `https://data.gov.uk/`

MHCLG, 2019. Total number of dwellings and net additional dwellings, borough.
URL `https://data.london.gov.uk/`

Monaco, F., Sali, G., Mazzocchi, C., Corsi, S., Jun 2016. Optimizing agricultural land use options for complying with food demand: evidences from linear programming in a metropolitan area. Aestimum 68, 45–59.

Nidumolu, U. B., van Keulen, H., Lubbers, M., Mapfumo, A., Jan 2007. Combining interactive multiple goal linear programming with an inter-stakeholder communication matrix to generate land use options. Environmental Modelling & Software 22 (1), 73–83.

ONS, 2017. Jobs in London by industry, 1996 to 2017.
URL `https://www.ons.gov.uk`

QGIS Development Team, 2019. QGIS Geographic Information System.
URL `http://qgis.osgeo.org`

Schlager, K., J., May 1965. A land use plan design model. Journal of the American Institute of Planners 31 (2), 103–111.