Proceedings of CONF-BPS 2026 Symposium: Innovation, Finance, and Governance for Sustainable Global Growth
DOI: 10.54254/2754-1169/2026.BJ31880

Designing Non-Price Allocation: A Market Design Analysis of NHS Healthcare Distribution

Zhengqi Ye

School of Humanities, University of Southampton, Southampton, UK
Zy4g24@soton.ac.uk

Abstract. This essay examines how scarce healthcare resources are allocated within the National Health Service and argues that the current framework fails to achieve adequate efficiency or fairness. Although NHS funding aims to secure equal access for equal need, the mechanisms that govern individual access,such as waiting-time rules, priority categories, and cost-effectiveness assessments, often operate at cross-purposes. Drawing on insights from mechanism design, this essay demonstrates how these rules resemble well-documented allocation challenges in non-price environments, where rigid priorities and overly broad classifications reliably generate inefficiencies. Ethical concerns reinforce these structural issues. QALY (Quality-Adjusted Life Year)-based assessments risk disadvantaging already vulnerable groups, while priority structures commonly fail to reflect morally relevant differences in severity. Combining these economic and ethical perspectives, the essay evaluates how existing NHS mechanisms distort access and why they struggle to align with the moral urgency of healthcare need. It then proposes reforms grounded in mechanism design, including fine-grained priority scoring, welfare-sensitive tie-breaking, flexible allocation rules, and adjustments to regional capacity. The analysis shows that a fair healthcare system cannot rely on fragmented procedures. Instead, it must be guided by coherent design principles that integrate both welfare consideration and moral judgment.

Keywords: Non-Price Allocation, Mechanism Design, Healthcare Resource Allocation, National Health Service (NHS), QALY and Priority Setting

1. Introduction

This essay examines how scarce healthcare resources are currently allocated within the National Health Service (NHS) and argue that the existing allocation framework is neither sufficiently efficient nor ethically justifiable. Although the NHS aims to secure "equal opportunity of access for equal need" through a population-based funding formula [1], the mechanisms that determine how patients actually receive treatment—particularly waiting-time rules, priority categories, and cost-effectiveness criteria, which generates systematic inequities. These mechanisms operate in a non-price environment where allocation is determined by institutional rules rather than market signals. Although Roth does not address the NHS explicitly, he argues that "market design… calls for an engineering approach", and the NHS allocation challenge aligns with this design-economics perspective.

DOI: 10.54254/2754-1169/2026.BJ31880

To assess the NHS framework, this essay draws on concepts from matching theory and mechanism design, focusing on priority structures, serial dictatorship-style waiting lists, and welfare-relevant inefficiencies. These economic concerns parallel core ethical objections in the healthcare rationing literature. As Harris argues, allocation systems that systematically disadvantage particular groups, such as the elderly, the disabled, or those with severe conditions, fail to treat individuals with equal moral concern [2]. Parfit's priority view similarly emphasises that benefiting the worse-off carries special moral significance [3]. Such considerations reveal that efficiency metrics alone cannot guide just allocation.

The remainder of this essay evaluates the NHS allocation mechanisms through this combined economic and ethical lens. Part I analyses how current NHS rule's structure actual patient access, highlighting the limitations of waiting-time systems and QALY-based assessments. Part II proposes mechanism-design reforms intended to improve efficiency while aligning allocation with the moral imperative of equal treatment.

2. NHS resource allocation and its structural limitations

The allocation of healthcare resources in the NHS is shaped by a combination of macro-level funding rules and micro-level access mechanisms. At the funding level, resources are distributed to regions according to a population-based formula designed to secure "equal opportunity of access for equal need" [1]. This formula incorporates age-weighted population measures, morbidity indices, and health inequalities adjustments, with the aim of aligning expenditure with underlying patterns of need [1]. While this approach reflects a commitment to fairness across regions, it does not directly determine how individual patients access care. The transition from population funding to individual access is mediated by institutional mechanisms—waiting lists, urgency classifications, and cost-effectiveness thresholds—which introduce new forms of structural disadvantage.

A central mechanism governing access to elective and specialist care is the waiting-time system. As the OECD's comparative study shows, waiting times emerge not merely from insufficient capacity but from the interaction of demand, supply, and administrative rules [4]. Under NHS rules, patients are typically prioritised according to a mixture of clinical urgency and time spent waiting. This resembles what matching theorists describe as a form of serial dictatorship: patients are treated roughly in the order they join the queue, unless they qualify for higher-priority categories. Serial dictatorship offers transparency but rarely achieves efficiency, as adhering to priority constraints can preclude Pareto-improving reallocations [5]. In practice, patients with moderate urgency may wait for longer than those whose expected benefit from treatment is comparatively small. The resulting allocation fails to track morally relevant differences between healthcare needs.

The fairness of waiting-time systems is further complicated by the incentives they create for providers. As Siciliani et al. note, waiting-time targets can distort provider behaviour. Resources are often redirected to cases that help meet performance metrics, rather than those with the greatest clinical need [4]. This reflects a form of strategic behaviour equivalent to what mechanism design theorists identify in school choice systems: when priorities and performance rules are poorly aligned with welfare, institutions may manipulate capacities or classifications to satisfy administrative requirements [5]. Such distortions undermine the normative foundations of NHS allocation, as procedural objectives are prioritised over substantive healthcare needs.

A further complication arises from the use of cost-effectiveness assessments in determining access to certain treatments. Drawing on QALY-based evaluations, NICE appraisals guide decisions on public funding for specific interventions. However, QALYs can systematically disadvantage individuals with limited baseline quality of life or shorter expected remaining lifespans—groups that

DOI: 10.54254/2754-1169/2026.BJ31880

already face structural disadvantage. As Harris argues, QALY-based frameworks can impose a form of double jeopardy: disabled patients are first disadvantaged by the lower quality-of-life scores assigned to their everyday experiences, and again when these scores reduce their priority for treatment [2]. If allocation principles deprioritise individuals who are already worse off simply because of their disadvantaged status, they fail to reflect the moral requirement of equal concern for all persons.

Parfit's priority view underscores this point. He maintains that morally, benefiting individuals carries greater weight the worse off those individuals are [3]. Yet QALY arithmetic, when used as a rationing principle, prioritize interventions that yield more life-years over those benefiting the most disadvantaged groups. It erases the distinction between saving persons and saving person-years. These are two moral categories that are not interchangeable [2]. This highlights a structural misalignment between efficiency metrics and ethical judgment. Mechanisms that prioritise efficiency in this reductive sense cannot reliably assess the urgency or severity of a patient's condition, nor can they recognize the moral significance of assisting disadvantaged individuals.

Collectively, these features of the NHS allocation system reveal a deeper structural flaw. Allocation rules, including waiting-time procedures, priority categories, and QALY-based assessments, operate independently and often conflict in their normative implications. Waiting lists benefit those who enrol earlier, and clinical urgency rules prioritize those with more rapidly deteriorating conditions, while QALY assessments favour those who can generate the greatest health gains. Each mechanism captures only one dimension of patients' moral needs. The result is a system that is neither consistently fair nor reliably efficient. Roth argues that institutions that allocate scarce resources without prices require coherent design principles to ensure that outcomes reflect legitimate social objectives [6]. Without such design, even well-intentioned rules can interact to produce outcomes that fail both morally and economically.

These considerations lay the groundwork for the reforms proposed in Part II. To allocate care fairly while enhancing overall welfare, the NHS must adopt mechanisms that integrate need, severity, and moral urgency. It should not rely on fragmented rules that obscure structural inequities

3. Mechanism-design-based reforms for fair and efficient NHS allocation

The limitations discussed in Part I demonstrate that the NHS needs a clearer and more consistent approach to care allocation, which better aligns its rules with practical welfare concerns and basic moral expectations. Mechanism design offers practical tools for achieving this aim. Since the NHS allocates scarce and indivisible resources without prices, its allocation challenges resemble classic one-sided matching environments. In such settings, the design of priority structures, tie-breaking rules, and exchange mechanisms directly influences outcomes. Roth describes this task as an engineering problem which institutions must be intentionally designed to produce outcomes that meet normative and welfare criteria [6]. Below are several mechanism design-based reforms that could improve fairness and efficiency while addressing the ethical concerns discussed earlier.

The first reform involves the refinement of priority structures. Current NHS priority categories, such as "urgent", "soon", and "routine", are coarse and fail to distinguish between morally significant differences in severity. In mechanism design terms, these categories act as rigid priority constraints that prevent welfare-improving [5]. A more fine-grained system would integrate medical severity scores, risk of deterioration, and expected intervention benefits. Such a structure would not merely increase efficiency but also better align with Parfit's priority view, which emphasizes aiding the worse-off [3]. Assigning proportionately higher priority to individuals facing the greatest disadvantage would ensure the allocation system better upholds the moral requirement of equal

DOI: 10.54254/2754-1169/2026.BJ31880

concern for all persons. It would prevent QALY arithmetic from dominating decisions with its bias towards those who can gain more life-years.

The second reform focuses on enhancing tie-breaking rules for patients within the same priority category. As discussed earlier, coarse categories generate numerous ties that require administrative resolution. In matching markets, arbitrary tie-breaking can lead to significant welfare losses by preventing matches that increase overall benefit [5]. The NHS could instead adopt tie-breaking rules that consider factors such as severity, moral urgency, and risk of irreversible harm. For example, when two "urgent" patients require treatment, their ordering should reflect differences in severity rather than procedural chance. This approach would address Harris's concern that existing allocation rules may exacerbate the disadvantage of already vulnerable individuals [2]. Incorporating morally salient features into tie-breaking, the NHS would ensure that administrative equality, which treats all "urgent" patients identically, does not undermine substantive moral equality.

The third reform draws on the top trading cycles (TTC) mechanism developed by Roth, Sönmez, and Ünver in the context of kidney exchange [7] TTC illustrates that even with scarce, indivisible resources, appropriately designed exchange mechanisms can enhance both fairness and efficiency. The key insight is that allocation rules can exploit complementarities between different agents' needs and capacities, enabling mutually beneficial exchanges that would otherwise be unachievable. Although elective-care allocation differs from kidney exchange, the underlying principle remains relevant: priority structures should permit welfare-improving reallocations when such actions do not violate legitimate claims. This idea challenges the rigidity of current waiting-time rules, which often prevent patients from receiving treatment sooner even when capacity exists and no other patient's rights would be infringed. Incorporating TTC-like flexibility, while preserving essential constraints, would enable the NHS to reduce inefficiencies without undermining equity.

The fourth reform concerns the structure of waiting lists themselves. As argued in Part I, NHS waiting lists resemble serial dictatorship: patients are admitted in the order of referral unless clinical urgency intervenes. While easy to administer, serial dictatorship is not universally efficient and overlooks relevant welfare information [5]. A more suitable alternative would be a deferred acceptance (DA)–style mechanism that ranks patients using a composite score reflecting severity, urgency, and expected benefit. Under DA, patients would be offered available slots in an order that more accurately reflects legitimate healthcare priorities. This would not exclude waiting time as a factor but would prevent time from overriding more morally salient considerations. Moreover, well-designed allocation mechanisms can reduce incentives for strategic behaviour by linking access to clear, predictable rules rather than performance-driven targets. This feature is central to numerous mechanism-design models, although Roth does not explicitly address hospitals or healthcare directly [6].

The final area for reform addresses regional inequalities in access. The NHS allocation formula aims to equalise opportunity by distributing funds based on need [1], yet significant regional disparities in waiting times remain. From a matching perspective, these inequalities resemble quota imbalances in many-to-one matching markets. If certain regions systematically receive more demand than they can accommodate, while others operate below capacity, overall welfare falls. A coordinated national mechanism for redistributing patients, particularly for high-severity cases, could reduce these disparities. Instead of treating regional boundaries as fixed constraints, the NHS could treat them as flexible quotas, ensuring individuals with the highest need receive treatment regardless of postcode. This reform aligns with Parfit's view that benefiting the worse-off has independent moral significance [3]. Regional identity is morally arbitrary; need is not.

DOI: 10.54254/2754-1169/2026.BJ31880

Together, these reforms illustrate how economic design and ethical reasoning can be mutually reinforcing. Mechanism design clarifies how institutional rules generate inefficiencies and inequities, while ethical analysis articulates why these outcomes are morally objectionable. Harris's critique of QALYs highlights the injustice of disadvantaging disabled patients or those with limited life expectancy, while Parfit's priority view provides a principled basis for prioritizing the severely ill. Mechanisms that integrate these considerations—fine-grained priority scores, welfare-sensitive tie-breaking, flexible exchange rules, and quota adjustments—would produce allocations that reflect both economic and moral judgment. They do not resolve the scarcity of healthcare resources, but they ensure scarcity is managed in a way that respects individuals as equals.

4. Conclusion

The analysis presented in this essay demonstrates that the NHS allocation framework, though guided by principles of equal access, is constrained by mechanisms that fail to capture morally relevant distinctions between patients. Waiting-time rules, coarse priority categories, and QALY-based assessments operate independently and produce outcomes that are neither consistently fair nor reliably efficient. Allocation without price mechanisms depends on institutional rules, and whether these rules work together coherently determines how fairly resources are distributed. When different procedures operate independently, well-intended rules can interact to obscure clinical need and reinforce existing inequalities.

Mechanism design offers practical tools for reform. More granular priority structures, welfare-sensitive tie-breaking, flexible allocation mechanisms inspired by TTC, and adjustments to regional quotas can enhance fairness and reduce inefficiency. Ethical considerations reinforce these reforms. Harris's critique shows that allocation principles must not systematically disadvantage the already vulnerable, while Parfit's priority view emphasizes the moral importance of addressing severe need. An allocation framework integrating these insights will better reflect the dual aims of efficiency and justice.

A just healthcare system should not simply maximise life-years, nor rely on rigid rules that disregard moral urgency. Instead, it should treat individuals as equals, respond to severity, and allocate resources through mechanisms that integrate both economic and ethical considerations.

References

[1] NHS England. (2023). Technical guide to allocation formulae and convergence revenue allocations: 2023/24 to 2024/25.
[2] Harris, J. (1987). QALYfying the value of life. Journal of Medical Ethics, 13(3), 117-123.
[3] Parfit, D. (1997). Equality and priority. Ratio, 10(3), 202-221.
[4] Siciliani, L., Borowitz, M., & Moran, V. (2013). Waiting time policies in the health sector: What works? OECD Publishing.
[5] Abdulkadiroğlu, A., & Sönmez, T. (2003). School choice: A mechanism design approach. American Economic Review, 93(3), 729-747.
[6] Roth, A. E. (2002). The economist as engineer: Game theory, experimentation, and computation as tools for design economics. Econometrica, 70(4), 1341–1378.
[7] Roth, A. E., Sönmez, T., & Ünver, M. U. (2004). Kidney exchange. Quarterly Journal of Economics, 119(2), 457-488.