

Effect of Resources on the Efficiency-Fairness Tradeoff for Allocation Problems

Shubham Singh¹, Chris Kanich¹, Ian A. Kash¹

¹University of Illinois Chicago

{ssing57, ckanich, iankash}@uic.edu

Abstract

Efficiency and fairness are conflicting objectives for resource allocation problems. Although well-studied, its ubiquity in numerous applications keeps presenting researchers with new challenges. In this work, we study the relationship of resources with efficiency and fairness and how it affects their tradeoff. We motivate the effect of resources on group utilities and the evaluation measures through the framework of homogeneous functions. Then, we explore the tradeoff together with three decision choices — allocating resources based on needs vs. outcomes, measuring social welfare as maximizing vs. minimizing objectives, and expressing fairness in terms of absolute vs. relative measures. Each choice presents us with a rich, complex setting and different desirable solutions. We illustrate our findings through stylized examples and then a more realistic example based on restaurant inspections.

1 Introduction

Resource allocation with respect to efficiency and fairness is often studied under the assumption of a fixed pool of resources. However, organizations often need to make a case for the benefits of being supplied with additional resources (e.g., when applying for grants). Similarly, a government facing budget cuts will care about the implications of reducing the resources provided to various departments. In such situations, it is not immediately clear *how the change in resources impacts efficiency and fairness*, since the change can affect efficiency and fairness not only individually but also together. For example, imagine a local government applying for a federal grant to improve public transit — the funds could be used to improve service (efficiency), start new lines that improve access (fairness), or improve both service and access (efficiency and fairness). Our paper approaches this problem by studying the change in the tradeoff between efficiency and fairness with a change in resources.

One might imagine that the effects of a change in resources could be understood simply by using any standard approach to the trade-off between efficiency and fairness at the new level of resources. As one example, we could study what would happen if we always chose the most efficient way to

deploy the resources and compare it to the fairest way as is done, e.g., in the literature on the Price of Fairness [Bertsimas *et al.*, 2011; Donahue and Kleinberg, 2020] (see Appendix A). Perhaps surprisingly, we find that such comparisons depend on subtle details of the way fairness and efficiency are measured. For example, given the *exact same* scenario different measurements will conclude that fairness increases linearly, sublinearly, or not at all with efficiency as resources are added. To study this interaction between resources and objectives, we zoom into the specific context to understand (i) how the tradeoff changes with a change in the amount of resources, (ii) how different definitions of societal objectives — efficiency and fairness — change the observed effect of resources.

On a technical level, we address (i) by explaining the effect of changing resources through the lens of homogeneity, a property of some functions that determines how they behave when scaling up their inputs. We observe that many common choices for utility functions and fairness measures are homogeneous and demonstrate how the degree of homogeneity affects the nature of tradeoffs. We use homogeneity to make several observations about tradeoff analysis. These observations address (ii) by highlighting how seemingly small technical details can lead to substantially different quantifications of tradeoffs and, thus, different decisions. In particular, we focus on the following three decisions:

- Should we measure utility in terms of the needs fulfilled or the outcomes created by fulfilling that need?
- Should we frame the optimization in terms of maximizing benefits or minimizing a cost function?
- Should we measure fairness violations in terms of the absolute or relative disparity in utilities?

All of these questions capture some aspect of the fact that efficiency and fairness are not dimensionless measures and none of these questions has an obvious right answer. Most combinations of answers seem reasonable for at least some plausible situation. But we expose how these technical decisions have important implications.

Need vs. outcomes captures the difference between analyzing the number of resources deployed to, e.g., each region of a city (adjusted based on population and other factors relevant to determining “need”) and what positive benefits the resources produce. From a mathematical perspective, the need

version is linear: each additional resource produces a constant benefit; measuring output is more complex but allows us to capture non-linearities: going from one street sweeper to two halves the time between cleanings, but going from two to three reduces it by a third. We show through a simple stylized example that the two formulations have very different implications for the fairness benefits of adding resources. In the former case adding resources will tend to exacerbate unfairness, as the efficient solution will concentrate them in the neediest region, while in the latter case the diminishing marginal returns of resources tend to lead to improved fairness when adding resources. Our point is not that one of these two choices is “right” (that may depend on what better captures the preferences of stakeholders in the situation), but rather that this simple modeling choice can substantially affect the way adding resources affects the tradeoff between efficiency and fairness.

Similarly, we could equally well define a utility function to capture the benefits of allocating resources or a societal cost function to capture the cost of only having allocated a limited number. (In the street sweeping example, this cost function might capture the amount of time the streets remain uncleaned.) This question turns out to have a non-obvious linkage to the third. We argue that the combination of measuring the utility of resources in terms of the reduction in societal costs and using relative disparity to measure fairness often leads to unreasonable conclusions about tradeoffs. In particular, it suggests that even with ample resources, so that there is little societal cost to any group, it is still possible for the efficient allocation to be extremely unfair. Returning to our earlier example, consider a city with two regions, region A and region B, receives a huge grant to hire more street sweepers such that it can reduce the overall street sweeping times. In region A, the sweeping time reduces from 2 days to 6 hours, and in region B, from 1 day to 3 hours. Although more resources are allocated to region A, we still reduce the disparity when measured in absolute terms and the overall time for which streets stay unclean. In relative terms, such as the ratio of sweeping times between the regions, the disparity indicates no improvement. We argue this is unreasonable; at some point, the gap will close enough that most people will agree that the outcome seems reasonably fair in practice, even if some difference remains.¹

In contrast, when a simple transformation of the efficiency axis reframes the optimization in terms of benefit, the *exact same scenario* would be evaluated as near-perfect fairness at even low resources. That is, if there are 30 days in a month then at the initial level of resources region A has 28 “clean days” while region B has 29, a ratio much closer to perfect fairness. This issue disappears when using absolute disparity to quantify fairness. At first, it seems absolute disparity presents us with the best choice possible, but it is less interpretable than relative fairness, and when utilities are treated as a proxy for needs (rather than outcomes), more resources decrease the fairness. Such caveats complicate making a rec-

ommendation in isolation and desire a more nuanced consideration. We propose that the choice of these different perspectives with changing resources can be understood by examining the homogeneity (or lack thereof) in the relevant definitions. In particular, homogeneity highlights the avoidable pitfalls of combining relative disparity and societal costs when resources can change.

Having illustrated these factors using a stylized model, we then show their relevance in a more realistic decision-making scenario with discrete, indivisible resources by examining their effects on the optimization of restaurant inspection schedules based on data from Chicago [Schenk Jr., 2015; Singh *et al.*, 2022]. We observe behavior broadly consistent with our homogeneity-based predictions.

In summary, our main contributions are:

- Proposing using the homogeneity of many common utility functions and objectives as a tool for understanding the effects of adding resources on the efficiency-fairness tradeoff. (§3 & 4)
- Demonstrating through stylized examples that three technical choices can have a substantial effect on the shape of these tradeoffs and that homogeneity can help explain this. (§5 & 6)
- Confirming that these insights are relevant in a scenario based on a real dataset that does not perfectly satisfy our stylized assumptions. (§7)

2 Related Work

Resource Allocation

Resource allocation has been studied in several forms and disciplines. Fair division literature has offered some notions of fairness that have been implemented in several applications. For example, Ghodsi *et al.* [2011] uses dominant resource fairness for scheduling computing tasks. For heterogeneous resources, researchers have used envy-based fairness notions [Wang *et al.*, 2024] and extended them to group fairness [Aleksandrov and Walsh, 2018]. Another prominent area for resource allocation is policy-making [Rodolfa *et al.*; Bastani *et al.*, 2021]. Several works have proposed simulation studies to determine fair resource allocation for government agencies [Rodolfa *et al.*; Nguyen *et al.*, 2023; Singh *et al.*, 2022], long-term fairness goals [Deb, 2011], COVID-19 testing [Bastani *et al.*, 2021], kidney exchange [McElfresh and Dickerson, 2018], and under censored feedback [Elzayn *et al.*, 2019]. Recent work has also studied counterfactual utilities for resource allocation [Kube *et al.*, 2023; Blandin and Kash, 2023]. Our work is inspired by these applications and studies the relationship between group utilities and scalable resources that determine the efficiency and fairness of a resource allocation solution.

Tradeoffs

The tradeoff between efficiency and fairness is intrinsic to decision-making with multiple objectives. Prior works have analyzed the tradeoffs between multiple fairness definitions [Kleinberg *et al.*, 2017; Friedler *et al.*, 2021; Her- tweek and R  z] and with performance and fairness measures [Brandao, 2021; Chen and Hooker, 2020; Liu *et al.*,

¹If six versus three hours still seems substantially unfair to you, what about six versus three minutes? Seconds? At some point, the claim that fairness has not improved becomes untenable.

2022; Radovanovic *et al.*, 2020; Andrews *et al.*, 2022; Heidari *et al.*, 2018; Singh *et al.*, 2022]. Recent work by Mashiat *et al.* [2022] presents resource allocation settings closest to ours with housing services that have non-uniform utility and capacities. They propose four fairness definitions that can be written as differences and ratios. Their findings state that same allocation can portray different groups as favored based on the chosen definition. Unlike their definition of tradeoff, which only considers fairness, we study the tradeoff between efficiency and fairness and the effect of resources on it.

3 Utility as Homogeneous Functions

Consider T different types of resources have to be allocated to G groups for $g \in \{1, \dots, G\}$. The number of resources allocated for each type is n_t for $t \in \{1, \dots, T\}$. Although each resource type is different, the resources within each type are assumed to be uniform and interchangeable. Mashiat *et al.* [2022]’s work on housing services illustrates one such setting where multiple resources capture different types of services. However, our formulation also works for cases with a single type of resource, for example, the number of public libraries in a neighborhood. Each group g receives some amount of each resource type and derives a benefit from this, which we call *group utility* denoted as u_g . We define *total utility*, or *social welfare* as an aggregate of the group utilities, which is a measure of total benefit drawn by all the groups from the allocated resources. An interesting property of the group utilities we explore in this section is homogeneity. When the group utilities are homogeneous functions of the resources allocated to them, it implies that if the allocated resources to a group are scaled by a factor of s , the resulting group utility is scaled by the same factor raised to some power k .

Definition 1 (Homogeneous Function). *A function $f : \mathbb{R}_{\geq 0}^T \mapsto \mathbb{R}$ is (positive) homogeneous of degree k if $f(sX) = s^k f(X)$, where s is a (positive) scalar and X is a T -dimensional vector of (positive) real numbers.*

Homogeneous utility functions allow us to model varied effects of the change in resources using group-specific functions. Even with a single type of resource, we can capture a wide variety of forms for $u_g(r_g)$ including linear, monomials such as r_g^2 , and even inverse relationships such as $\frac{1}{r_g}$ if it is an undesirable thing such as polluting industrial activity. In a resource allocation setting, our objectives would combine (in the case of efficiency) or compare (in the case of fairness) the benefit resource units $\mathbf{r} \in \mathbb{R}_{\geq 0}^T$ bring to the groups. If we have an objective function f measuring fairness of efficiency that is also homogeneous, the resulting combination is homogeneous as well:

Proposition 1. *If the group utilities are homogeneous functions of degree k_1 , and the objective function is defined as a homogeneous function of degree k_2 over the group utilities, then the resulting combination is a homogeneous function of degree $k_1 k_2$.*

Proof. The statement is direct from the definition of homogeneity. Consider G groups whose utility functions are denoted as $u_g(\cdot)$, and let $\mathbf{r}_g = [r_1, \dots, r_T]$ denote the resource

allocation.

$$\begin{aligned} f(u_1(s\mathbf{r}_1), \dots, u_G(s\mathbf{r}_G)) &= f(s^{k_1} u_1(\mathbf{r}_1), \dots, s^{k_1} u_G(\mathbf{r}_G)) \\ &= s^{k_1 k_2} f(u_1(\mathbf{r}_1), \dots, u_G(\mathbf{r}_G)) \quad \square \end{aligned}$$

Using Proposition 1, we present two applications where the choice of the objective function can lead to different outcomes based on the change in group resources. First, consider a linear function of the group utilities as an objective. This is common when considering efficiency measures like social welfare and fairness measures that consider the difference in group utilities.² A linear function is homogeneous of degree 1. Thus, the effect of change in resources on linear functions is to scale the objective accordingly.

Corollary 1. *If the group utilities are homogeneous functions of degree k , then a linear function of the group utilities is also a homogeneous function of degree k ;*

$$f(u_1(s\mathbf{r}_1), \dots, u_G(s\mathbf{r}_G)) = s^k f(u_1(\mathbf{r}_1), \dots, u_G(\mathbf{r}_G)).$$

Second, we explore another natural choice for objective functions, a ratio of the group utilities. Ratios are particularly amenable when quantifying comparisons as they are relative – compared to a known quantity (denominator), independent of units, and could be expressed in fixed intervals, making them easy to understand. Since ratios are defined over two quantities, for this example, let us assume $g = 2$, but our results extend to any number of g via simple transformations, for example, comparing extrema over a set.³

Corollary 2. *If the group utilities are homogeneous functions of degree k , then their ratio is a homogeneous function of degree zero. Thus,*

$$\frac{u_1(s\mathbf{r}_1)}{u_2(s\mathbf{r}_2)} = \frac{u_1(\mathbf{r}_1)}{u_2(\mathbf{r}_2)}.$$

An important implication of this corollary is that $f(\cdot)$, when calculated as a ratio of the group utilities, the result is independent of any changes in the total resources $\mathbf{r}$. This result highlights an important caveat of using ratio as an objective function, and we further examine the implications of this behavior in the later sections.

Of course, these results only apply when the resources allocated to all groups are scaled up by the same constant. It is not immediately obvious that we should do so; perhaps it would be better to allocate the additional resources disproportionately to one group. However, for homogeneous functions, scaling up evenly is optimal.

Proposition 2. *Let $\mathbf{r}^*$ be the optimal allocation of some quantity of resources given group utilities which are homogeneous of degree k_1 and an objective f which is homogeneous of degree k_2 . Then given s times as many resources the optimal allocation is $s\mathbf{r}^*$.*

²Defining fairness in terms of group utilities is a way to extend standard group fairness notions beyond classification settings [Blandin and Kash, 2023], and has been used in prior work on resource allocation [Donahue and Kleinberg, 2020; Singh *et al.*, 2022; Globus-Harris *et al.*, 2022]

³Demographic Parity Ratio (DPR) for more than two groups is often defined as the worst case, min-max ratio [Ghosh *et al.*, 2021; Weerts *et al.*, 2023].




Figure 1: Hypothetical Efficiency-Fairness tradeoffs at varying amounts of resources. The circles joined by the dashed lines show one possible increase in fairness and efficiency obtained by increasing the resources. The other dashed lines show alternate possibilities.

Proof. Any other allocation can be written in the form $s\mathbf{r}$ where $\mathbf{r}^*$ and $\mathbf{r}$ allocate the same total amount of each resource. Then we have:

$$\begin{aligned} f(u_1(s\mathbf{r}_1^*), \dots, u_G(s\mathbf{r}_G^*)) &= s^{k_1 k_2} f(u_1(\mathbf{r}_1^*), \dots, u_G(\mathbf{r}_G^*)) \\ &\geq s^{k_1 k_2} f(u_1(\mathbf{r}_1), \dots, u_G(\mathbf{r}_G)) \\ &= f(u_1(s\mathbf{r}_1), \dots, u_G(s\mathbf{r}_G)), \end{aligned}$$

where the two equalities follows by Proposition 1 and the inequality follows from the optimality of $\mathbf{r}^*$. $\square$

4 Resources as New Dimension of Fairness

The total amount of resources is one of the largest constraints when designing the algorithm for a socio-technical problem. This is often presented as an immutable dimension and, therefore, often overlooked when considering the desirable efficiency and fairness properties of an algorithm. However, we argue that providing resources can be as important as the sacrifice of efficiency, making this dimension equally important when considering fairness tradeoffs.

Let's revisit the example of a local government applying for a federal grant (see §1) that makes more resources available. Naively applied, without changing any algorithm parameterization, more resources may increase efficiency but lower fairness or vice versa. Or more resources can simultaneously increase both efficiency and fairness. More broadly, we could retune our algorithm and the exact set of choices available to us would vary based on the upper and lower bounds of the solutions to the multi-objective optimization.

We show a hypothetical example of the transformation to the relationship between efficiency and fairness based on the resources available in Fig. 1. On the X-axis, we show a hypothetical homogeneous fairness function $F(\cdot)$, and on the Y-axis, a homogeneous efficiency function $E(\cdot)$. Each tradeoff curve shows the scaling of the efficiency and the fairness functions for 100, 200, 300 resources. Each curve presents a different Pareto front offered by scaling the resources and each individual marker shows one possible allocation along with its efficiency and fairness. Thus, this figure captures how the tradeoff curve changes with resources, in this case keeping its shape (as in Corollary 1).

Let's say our current amount of resources is 200 (orange triangles), and our current operating point is illustrated by

the circle on the tradeoff curve in Fig. 1. The dashed lines between the operating point associated with 200 resources and the one on the 300 curve (green crosses) show a possible improvement along both axes – efficiency (Y-axis) and fairness (X-axis), both individually and together, from a grant that added 100 resources. If we use the increase in resources purely to improve fairness, the resulting improvement is essentially the same as keeping the resources the same and moving fully to the fair end of the Pareto frontier. Similarly, the efficiency achieved by the increased resources (vertical dashed line) is beyond what is possible with 200 resources. Faced with this scenario, we would argue that the most important decision is likely not the tradeoff between efficiency and fairness but the tradeoffs involved in getting more resources. Thus, it is important to quantify this third dimension when examining such tradeoffs. Using this same example (Fig. 1) in Appendix A, we show how PoF can lead to incomplete conclusions about the tradeoffs and how homogeneity can explain this.

5 Utility — Need vs. Outcome

Imagine a scenario in which groups have different utilities and are known to us. We want to explore how the choice of measuring the efficiency based on the *needs* of those groups compares to the *outcomes* those groups desire. Such alternative ways to calculate efficiency can present decision-makers with different magnitudes of feasible, competing measures.

Consider two groups, $g \in \{g_1, g_2\}$, among which we need to divide the n resources of a common type $|T| = 1$, and values of n range from 2 to 10. Suppose we know the two groups have different utilities, $u_g(r_g)$, gained from receiving some amount of resource, r_g . We define their *need*-based utilities as:

$$u_{g_1}(r_{g_1}) = 2r_{g_1}, \quad u_{g_2}(r_{g_2}) = r_{g_2}. \quad (1)$$

Here, one group, g_1 , derives twice as much utility from receiving r resources as the other group, g_2 . This assumption might reflect, for example, greater importance of the resource to g_1 or a greater capacity to use it effectively. Let's consider two methods of resource allocation among the groups, $\{g_1, g_2\}$. An *efficient allocation* would assign all the resources to g_1 , given its higher utility. On the other hand, a *fair allocation* would split the resources among the groups so that each group gets equal utility from their allocation. Our goal is to measure the efficiency and fairness of the allocations. We compute the *efficiency* for the allocations by averaging the sum of utilities over the groups. We can calculate the *unfairness* by using the L^2 norm over the difference between each group's utility and efficiency of the allocation. As is often done, we transform the unfairness into a maximizing fairness function by subtracting the unfairness from a constant, C , so that the fairness values are non-negative. For our example, we chose C based on the highest value of unfairness achieved across levels of resources, but C could be any arbitrarily high value to ensure boundedness. Since we compute the L^2 norm of two absolute values, we term it *absolute fairness*. The efficiency and absolute fairness functions are given




Figure 2: The efficient allocation (blue circles) and the fair allocation (orange triangles) show the ends of the Pareto front representing an increasing number of resources. Efficiency and fairness are computed using *need*-based group utilities.

as follows:

$$E(\mathbf{r}) = \frac{1}{|\{g_1, g_2\}|} \sum_{g \in \{g_1, g_2\}} u_g(r_g), \quad (2)$$

$$F_{\text{abs}}(\mathbf{r}, E(\mathbf{r})) = C - \sqrt{\sum_{g \in \{g_1, g_2\}} (u_g(r_g) - E(\mathbf{r}))^2}. \quad (3)$$

Absolute fairness is a widely adopted definition of (un)fairness for multi-group settings [Dwork *et al.*, 2018; Donini *et al.*, 2018], although the choice of the norm may vary. We show the efficiency-fairness tradeoff curves in Fig. 2, where each pair of circles and crosses show the ends of a Pareto front corresponding to one choice of n . The lowest blue and orange marks correspond to 2, while the highest corresponds to 10. We note that the fairness (X-axis) for the fair allocation (orange triangles) stays the same since regardless of the number of resources, we can choose a perfectly fair solution. However, the increase in resources produces an increase in the efficiency of this fair allocation. In contrast, efficient allocation starts with lower absolute fairness, and as the resources are added, the fairness decreases as one group’s utility dominates the allocation. We can characterize the effect of resources for efficient allocation as a property of homogeneity (see §3). Since efficiency, fairness⁴, and both utilities are all homogeneous of degree 1, we observe both axes changing by a constant factor. As a result, the gap between the allocations with maximum fairness and maximum efficiency becomes wider as the resources increase. This example highlights the caveats of adding resources when they have a linear relationship with the group utilities and could be thought of *need*-based allocation. A real-world example of such a scenario would be allocating the number of police officers to neighborhoods by the crime-per-capita metric. If a decision-maker is not deliberate about how the resources are assigned in this setting, the results might be much worse along one axis than the other.

Now let’s consider another form of a utility function, *outcome*-based. Imagine each group has a fixed number of jobs that must be completed. In that case, adding an extra resource will speed up the completion time, increasing the group util-

⁴It is only homogeneous without the constant C . As we are interested in the difference between points, the C naturally cancels.

ity by the inverse factor. This yields a utility of the form:

$$u_{g_1}(r_{g_1}) = 1 - \frac{1}{2r_{g_1}}, \quad u_{g_2}(r_{g_2}) = 1 - \frac{1}{r_{g_2}}. \quad (4)$$

Analogous to *need*-based utility functions described in Eq. 1, one group, g_1 benefits more from the same amount of resources than g_2 . Using the same allocation mechanisms described above and the methods to calculate efficiency and absolute fairness, we plot the tradeoff curves for different resources in Fig. 3a. Unlike Fig. 2, the fairness (X-axis) of the efficient allocation increases as we add more resources, as increasing resources brings the utilities for both groups closer to 1. Because the efficient solution has now changed, the inverse relationship between the resources and group utilities causes the gap between the extreme points of the Pareto front – efficient allocation (blue circles) and fair allocation (orange triangles) to shrink with the addition of resources. Since the (differences in) group utilities are homogeneous of degree -1, changes in measures of efficiency and fairness diminish with added resources. In this type of scenario, where there are diminishing returns to providing additional resources, with enough resources we would naturally converge towards a more efficient solution for a low cost of fairness. A real-world example includes buying more snowplowing machines for city neighborhoods. Adding more machines helps all the regions and the effect of *fair* or *efficient* allocation starts to vanish after a while.

Our point is not that either formulation of utility is better than the other. It is up to stakeholders whether the importance of, e.g., police is in providing a feeling of safety through their presence or in their effectiveness at preventing crime. Similarly, fairness in one case corresponds to no group feeling over- or under-policed relative to others, while the other corresponds to equalizing the crime experienced. All of these may be reasonable goals. Our point is that choosing one or the other (or even a combination) can have a substantial impact on the nature of efficiency-fairness tradeoffs, and that homogeneity can help us understand the shape of impact we should expect.

6 Expression of Goals

In this section, we explore the example with the outcome-based utilities (Eq. 4) to investigate how different expressions of efficiency and fairness can modify their purpose and choice for the best decision. We modify both, together and separately, the fairness measure from absolute to relative and switch the interpretation of efficiency from maximizing a utility function to minimizing a cost function.

In addition to *absolute* fairness (Eq. 3), another expression for fairness is through ratios. *Relative* fairness could be defined as the geometric mean of the ratios of group utilities and allocation efficiency, given as:

$$F_{\text{rel}}(\mathbf{r}, E(\mathbf{r})) = \prod_{\{g_1, g_2\}} \left(\frac{\min(u_g(r_g), E(\mathbf{r}))}{\max(u_g(r_g), E(\mathbf{r}))} \right)^{\frac{1}{|\{g_1, g_2\}|}} \quad (5)$$

This relative expression of fairness is useful when the fairness goals are set based on the other group utilities and independent of the actual units. We show the efficiency-fairness




Figure 3: The Pareto front with an increasing number of resources when efficiency (Y-axis) and fairness (X-axis) are computed using *outcome*-based group utilities.

tradeoff and resources in Fig. 3b. Similar to Fig. 3a, the Pareto front contracts as we increase the resources, and the efficient allocation moves closer to the fair allocation. However, the relative fairness measure shows that efficient allocation moves more towards fair allocation when adding resources at lower resource levels than the absolute fairness measure. In other words, the improvement in relative fairness appears to be sublinear compared to the linear growth with absolute fairness. This is despite the efficient allocation being *exactly the same* in the two cases. Although both measures of fairness are not directly comparable, we can identify ways they lead to different conclusions. For example, the slope of the line between the efficient and fair solutions is steeper at low resource values and shallower at high resource values in Fig. 3b than Fig. 3a. This might lead to different decisions about how much utility to trade for fairness depending on the choice of absolute vs. relative.

Thus far, we have only looked at the efficiency measures that must be maximized for social good. Now, we explore the effect of a cost function that must be minimized. Continuing with our example of *outcome*-based utilities, we slightly change the group utilities keeping the privileged group as-is. They are given as follows:

$$c_{g_1}(r) = \frac{1}{r}, \quad c_{g_2}(r) = \frac{1}{2r}. \quad (6)$$

We compute the *societal cost* similar to the *efficiency* (Eq. 2), but the utilities now aggregate to a quantity whose objective is transformed into a reduction in loss incurred to the society when a resource is unavailable. From an optimization perspective, this transformation is harmless: maximizing efficiency and minimizing cost are equivalent. Combining the homogeneity of cost (degree -1), societal cost (degree 1), absolute fairness (degree 1), and relative fairness (degree 0), we can now apply our theoretical results to these scenarios.

We contrast the cost vs. absolute fairness and relative fairness in Figs. 3c & 3d respectively. Consistent with the findings in Fig. 3a, Fig. 3c shows the shape of the Pareto front inverted across a line parallel to the X-axis, and it shrinks as the cost for the efficient and fair allocation reduces. Proposition 1 tells us that both axes are homogeneous of degree -1 , which is why each doubling of the resources halves the gap to optimality on each. On the other hand, Fig. 3d illustrates an interesting result. When measuring relative fairness with costs, the fairness (X-axis) for efficient allocation (blue circles) and fair allocation (orange triangles) *do not change* with

the increase in resources, and the tradeoff stays constant. This is because the ratio of group costs remains the same for both the efficient and fair allocations and is a direct application of Corollary 2. Again, the underlying efficient allocations in the two figures are *exactly the same*. However, conclusions drawn from the two figures would differ vastly and could lead to different choices. For example, by only looking at Fig. 3d, we could infer that the addition of resources has no effect on the fairness of the efficient allocation. Conversely, measuring efficiency with relative fairness (Fig. 3b) shows the increase in fairness comes at a low cost of efficiency with rising resources, which we view as a more reasonable reflection of the situation.

The key takeaway is that seemingly small technical decisions, such as the choice to measure fairness violations in relative vs. absolute terms, can have substantial effects on the nature of the efficiency-fairness tradeoff with the change in resources and that these effects can be understood through the lens of homogeneity. In Appendix B, we perform a similar analysis for need-based utilities (linear so homogeneous of degree 1), and observe again that such technical choices have important consequences.

7 Empirical Study: Chicago Food Inspections

Now, we explore a real-world scenario where noise and other factors mean utility and objective functions do not precisely satisfy the definition of homogeneity. More broadly, this setting is much richer than our stylized examples (e.g. it has nine groups). Our findings from the empirical study of the Chicago food inspections dataset support our insights about the effect of resources on the efficiency-fairness tradeoff and choice of objectives, even when our assumptions hold imperfectly.

Food inspections are part of the services provided by local governments and involve agents (or sanitarians) visiting food establishments and evaluating the conditions as per food safety code [Chicago Department of Public Health, 2023]. The agents' objective is to identify and help correct any *critical violations* of the food code, as the critical violations pose the highest risk of producing conditions that might result in a foodborne illness. The primary dataset we explore contains more than 18,000 records of inspections done by the agents from the Chicago Department of Public Health (CDPH) over four years, from September 2011 to October 2014 and is pub-




Figure 4: Efficiency-Fairness tradeoffs correspond to various amounts of resources (shown in the legend) for the Chicago Food inspection dataset. Error bars represent standard error from 16-fold cross validation.

licly accessible on Chicago Data Portal⁵. Prior studies have used the dataset to train a machine learning model that prioritizes the inspections with a higher likelihood of resulting in critical violations [Schenk Jr., 2015; Kannan *et al.*, 2019; Singh *et al.*, 2022]. The model outputs a score that is used to order inspections from the highest score to the lowest to get a schedule that aims to find the critical violations early.

We describe our methodology in detail in Appendix C, including how we define utility and cost functions based on the outcome of time to detect critical violations, how we choose inspection schedules corresponding to efficient and fair allocations, and how we simulated creating these inspection schedules with varying numbers of agents available to conduct inspections. This allows us to create an analogue of Fig. 3 derived from a real application, shown in Fig. 4.

Comparing Fig. 3a to Fig. 4a, we see they are consistent in that the increase in resources makes the tradeoff between efficiency and fairness smaller. One key difference from our stylized example in Fig. 3a is that the fair allocation does not achieve the highest possible absolute fairness. This can be explained by the two-step scheduling process: first the ML model is used to get the likelihood score, and then the score is used to order the inspections. However, the model is not aware of the true utility of the scheduling steps and can not be optimized to achieve equal group utilities *ex-post*. Nevertheless, the figure still presents a strong argument for the addition of resources as a way to jointly improve fairness as the tradeoff between the two becomes smaller and the number of agents grows.

Next, we observe the change in tradeoffs when measuring the *relative* fairness instead of absolute fairness, shown in Fig. 4b. The plot reflects our observations from Fig. 3b, the efficient allocation gets closer to the fair allocation as the resources increase with a seemingly smaller benefit of choosing the fair solution under relative fairness than under absolute fairness. The caveat about the fair allocation being unable to achieve perfect fairness still shows up here. In contrast to absolute fairness, relative fairness might lead to the conclusion that even with a low number of resources, the system is already quite fair (note the different scales on the x-axes), and in general, choosing efficient rather than fair allocations

appears much more appealing.

We illustrate the effect of resources on a minimizing societal cost function along with both – absolute and relative fairness in Figs. 4c and 4d, respectively. Like the example in Fig. 3c, Fig. 4c shows the tradeoff to be shrinking as the resources are added for *efficient* and *fair* allocations. We note that the fair allocation does not achieve perfect fairness in this case as it is presented with the same challenge of finding out the true utilities *a priori* as when using the maximizing efficiency case (Figs. 4a & 4b). Fig. 4d echoes the findings and limitations of using cost and relative fairness together from Fig. 3d. As the resources increase, the cost of both the efficient and fair allocations decreases. But similar to Fig. 3d, and despite the added complexity of the real scenario, the Fig. 4d tradeoff largely does not change, as predicted by homogeneity. Once again, we could come to the wrong conclusions if we were to choose this combination of efficiency and fairness measures. In this application, both cost and relative fairness individually pose plausible choices, but their combination could lead to adverse decisions.

8 Conclusion

We have argued for considering the effect of resources on the efficiency-fairness tradeoff and provided several insights about the way modeling choices can influence decisions. In particular, we examined how the choices of defining utility functions in terms of need or outcome, defining fairness in absolute or relative terms, and representing goals as benefits or costs. We have seen in both stylized examples and a scenario derived from a real dataset that these technical modeling choices can lead to very different conclusions about the tradeoff between efficiency and fairness and the way it changes with the addition or removal of resources. Furthermore, these differences are predictable and can be understood through the lens of homogeneous functions. We do not provide reductive advice that there is a single “correct” choice: these decisions depend on details of the application domain and the values of those affected. But our work highlights the importance of thinking about them and provides a framework for understanding their consequences, particularly in the context of important societal decisions about resource allocation.

⁵<https://data.cityofchicago.org/>

References

- Martin Aleksandrov and Toby Walsh. Group Envy Freeness and Group Pareto Efficiency in Fair Division with Indivisible Items. In Frank Trollmann and Anni-Yasmin Turhan, editors, *KI 2018: Advances in Artificial Intelligence*, pages 57–72. Springer International Publishing, 2018.
- Kenya Andrews, Mesrob Ohannessian, and Tanya Berger-Wolf. Modeling Access Differences to Reduce Disparity in Resource Allocation. In *Equity and Access in Algorithms, Mechanisms, and Optimization*, pages 1–11, Arlington VA USA, October 2022. ACM.
- Hamsa Bastani, Kimon Drakopoulos, Vishal Gupta, Ioannis Vlachogiannis, Christos Hadjicristodoulou, Pagona Lagiou, Gkikas Magiorkinis, Dimitrios Paraskevis, and Sotirios Tsiodras. Efficient and targeted COVID-19 border testing via reinforcement learning. *Nature*, 599(7883):108–113, November 2021.
- Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. The Price of Fairness. *Operations Research*, 59(1):17–31, February 2011.
- Jack Blandin and Ian A Kash. Generalizing group fairness in machine learning via utilities. *Journal of Artificial Intelligence Research*, 78:747–780, 2023.
- Martim Brandao. Socially Fair Coverage: The Fairness Problem in Coverage Planning and a New Anytime-Fair Method. In *2021 IEEE International Conference on Advanced Robotics and Its Social Impacts (ARSO)*, pages 227–233, Tokoname, Japan, July 2021. IEEE.
- Violet (Xinying) Chen and J. N. Hooker. A Just Approach Balancing Rawlsian Leximax Fairness and Utilitarianism. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’20, pages 221–227, New York, NY, USA, February 2020. Association for Computing Machinery.
- Chicago Department of Public Health. City of Chicago: Food code rules. https://www.chicago.gov/content/dam/city/depts/cdph/FoodProtection/2023/Food-Code-Rules-Proposed-Revisions_05.31.23.pdf, 2023. Accessed Aug 9, 2024.
- Kalyanmoy Deb. Multi-Objective Optimization Using Evolutionary Algorithms: An Introduction. February 2011.
- Kate Donahue and Jon Kleinberg. Fairness and utilization in allocating resources with uncertain demand. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 658–668, Barcelona Spain, January 2020. ACM.
- Michele Donini, Luca Oneto, Shai Ben-David, John S Shawe-Taylor, and Massimiliano Pontil. Empirical Risk Minimization Under Fairness Constraints. page 11, 2018.
- Cynthia Dwork, Nicole Immorlica, Adam Tauman Kalai, and Max Leiserson. Decoupled Classifiers for Group-Fair and Efficient Machine Learning. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, pages 119–133. PMLR, January 2018.
- Hadi Elzayn, Shahin Jabbari, Christopher Jung, Michael Kearns, Seth Neel, Aaron Roth, and Zachary Schutzman. Fair Algorithms for Learning in Allocation Problems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 170–179, Atlanta GA USA, January 2019. ACM.
- Sorelle A. Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. The (Im)possibility of fairness: Different value systems require different mechanisms for fair decision making. *Communications of the ACM*, 64(4):136–143, April 2021.
- Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. Dominant Resource Fairness: Fair Allocation of Multiple Resource Types. 2011.
- Avijit Ghosh, Lea Genuit, and Mary Reagan. Characterizing Intersectional Group Fairness with Worst-Case Comparisons. In *Proceedings of 2nd Workshop on Diversity in Artificial Intelligence (AIDBEI)*, pages 22–34. PMLR, 2021.
- Ira Globus-Harris, Michael Kearns, and Aaron Roth. An Algorithmic Framework for Bias Bounties. 2022.
- Hoda Heidari, Claudio Ferrari, Krishna Gummadi, and Andreas Krause. Fairness Behind a Veil of Ignorance: A Welfare Analysis for Automated Decision Making. 2018.
- Corinna Hertweck and Tim R  z. Gradual (In)Compatibility of Fairness Criteria. 36(11):11926–11934.
- Vinesh Kannan, Matthew A. Shapiro, and Mustafa Bilgic. Hindsight Analysis of the Chicago Food Inspection Forecasting Model. *arXiv:1910.04906 [cs, stat]*, October 2019.
- Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent Trade-Offs in the Fair Determination of Risk Scores. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, volume 67 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 43:1–43:23, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- Amanda R. Kube, Sanmay Das, and Patrick J. Fowler. Fair and Efficient Allocation of Scarce Resources Based on Predicted Outcomes: Implications for Homeless Service Delivery. 76:1219–1245, 2023.
- Tennison Liu, Alex J. Chan, Boris van Breugel, and Mihaela van der Schaar. Practical Approaches for Fair Learning with Multitype and Multivariate Sensitive Attributes, November 2022.
- Tasfia Mashiat, Xavier Gitiaux, Huzefa Rangwala, Patrick Fowler, and Sanmay Das. Trade-offs between Group Fairness Metrics in Societal Resource Allocation. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 1095–1105. ACM, 2022.
- Duncan McElfresh and John Dickerson. Balancing Lexicographic Fairness and a Utilitarian Objective With Application to Kidney Exchange. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), April 2018.

- Khanh Duy Nguyen, Nima Shahbazi, and Abolfazl Asudeh. PopSim: An Individual-level Population Simulator for Equitable Allocation of City Resources, April 2023.
- Sandro Radovanovic, Andrija Petrovic, Boris Delibasic, and Milija Suknovic. Enforcing fairness in logistic regression algorithm. In *2020 International Conference on INnovations in Intelligent SysTems and Applications (INISTA)*, pages 1–7, Novi Sad, Serbia, August 2020. IEEE.
- Kit T. Rodolfa, Hemank Lamba, and Rayid Ghani. Empirical observation of negligible fairness–accuracy trade-offs in machine learning for public policy. *3*(10):896–904.
- Tom Schenk Jr. Forecasting restaurants with critical violations in Chicago. Technical report, May 2015.
- Shubham Singh, Bhuvni Shah, Chris Kanich, and Ian A. Kash. Fair decision-making for food inspections. In *Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO ’22, New York, NY, USA, 2022. Association for Computing Machinery.
- Haibin Wang, Sujoy Sikdar, Xiaoxi Guo, Lirong Xia, Yongzhi Cao, and Hanpin Wang. Multi-Type Resource Allocation with Partial Preferences. *34*(02):2260–2267, 2024.
- Hilde Weerts, Miroslav Dudík, Richard Edgar, Adrin Jalali, Roman Lutz, and Michael Madaio. Fairlearn: Assessing and improving fairness of ai systems, 2023.

A Effect of Resources on Price of Fairness

Quantifying the tradeoff between efficiency and fairness is a well-studied problem. We have shown the tradeoff as the Pareto fronts over maximizing and minimizing efficiency functions and *absolute* and *relative* fairness functions. Another popular albeit concise representation of this tradeoff is the *Price of Fairness (PoF)* [Bertsimas *et al.*, 2011; Donahue and Kleinberg, 2020]. PoF is defined⁶ as the ratio of the maximum utility from the sets of all solutions S to the maximum utility from the set of fair solutions $F \subseteq S$, given as:

$$\text{PoF} = \frac{\max_{s \in S} \text{efficiency}(s)}{\max_{s \in F} \text{efficiency}(s)} \quad (7)$$

Note this definition computes the ratio of ‘efficiency’ of the most efficient solution and the fairest solution to specify the tradeoff. Applied to the example we present in Fig. 1, the definition could be rewritten as:

$$\text{PoF} = \frac{\max(E(r))}{\max_{r \in \arg \max F(r)} E(r)} \quad (8)$$

By looking at Fig. 1, we can say that the choice of picking a certain amount of efficiency to trade for fairness would have vastly different consequences for the stakeholders in the three scenarios. However, if we just considered the PoF for the curves, it is uninformative since the ratio of the highest point to the lowest point on the Y-axis (efficiency) does not change as $E(r)$ is scaled. Therefore, if the choice were to be made solely by looking at the PoF for the three curves, the decision maker would not see the difference between the three tradeoff curves.⁷

The invariance of PoF for any change in resources can easily be explained by considering the homogeneity of the function. As stated in Corollary 2, ratios are homogeneous functions of degree 0 and, therefore, do not capture the effect of scaling by a constant factor. This example presents a setting where, while the PoF provides useful information about the implications of a commitment to fairness, it alone does not portray the complete picture necessary for comparing efficiency-fairness-resources tradeoffs.

B Tradeoff Measures for Linear Utilities

In §6, we explore how the expression of efficiency and fairness with *outcome*-based (inverse) group utilities changes the efficiency-fairness-tradeoff and, consequently, the inferences from the allocation methods. In this section, we present a similar analysis but with *need*-based (linear) utilities.

Fig. 5 presents the tradeoff for efficiency, fairness, and resources when using efficiency and relative fairness. Although the efficiency of both allocation methods increases with the




Figure 5: The three-way tradeoff curves for *need*-based (linear) group utilities measuring efficiency (Y-axis) and relative fairness (X-axis). Higher values for both are better.




Figure 6: The tradeoff curves for linear group utilities considering cost (Y-axis) and absolute fairness (X-axis). Lower cost and higher absolute fairness values are better.

increase in resources, it does not affect the fairness of efficient and fair allocation. Unlike Fig. 3d, this is an accurate reflection of the situation: no matter how many resources we have, giving them all to one group (the efficient solution) is perfectly unfair. This example is somewhat degenerate in this regard.

Now, to transform the linear utilities given in Eq. 1 into a minimizing cost function, we subtract the group utilities from a constant large enough to keep the utilities non-negative, and g_1 remains the privileged group. We write them as:

$$c_{g_1}(r) = \beta - r, \quad c_{g_2}(r) = \beta - 2r. \quad (9)$$

Using these utilities, we plot the Pareto front for cost and absolute fairness in Fig. 6. We notice a steady decline in the fairness measure, along with the reduction in cost for the efficient allocation with the addition of more resources. The fair allocation retains the highest absolute fairness but also a higher cost when compared to the efficient allocation at the same level of resources.

Finally, we show the tradeoff curve for cost vs. relative fairness in Fig. 7. Similar to Fig. 6, the ends of the Pareto front widen as the efficient allocation gets less fair and cost-effective, while the fair allocation stays perfectly fair and with a lower reduction in cost as the resources are added. However, relative fairness portrays that efficient allocation continues to have a higher level of fairness for fewer resources than when measuring absolute fairness in Fig. 6. Since linear costs are not homogeneous due to the constant β , Fig. 7 does not ex-

⁶The two papers actually use slightly different definitions, though similar in spirit. Interestingly, they differ in a way that can affect decisions. We use the latter.

⁷Donahue and Kleinberg [2020] more generally allow a bounded amount of unfairness, but similar examples can be constructed. The key point remains that while the PoF can give us some information about the Pareto frontier, it does not tell us how to choose a point on it.




Figure 7: The tradeoff curves for linear group utilities considering the cost (Y-axis) and relative fairness (X-axis). Lower cost and higher relative fairness values are desired.

hibit the problematic behavior we saw in Fig. 3d.

C Methodology for Empirical Study

Food inspections are one of the essential public services provided by a local jurisdiction to ensure public safety. Typically, food inspection involves agents (or sanitarians) who visit food establishments and check the food preparation environment and food handling practices against a food safety code [Chicago Department of Public Health, 2023]. When observed, the agents can cite the food establishments with violations, which are of three classes: *minor*, *serious*, and *critical*, with *critical* violations being the strongest signal of food resulting in an illness. Therefore, if an establishment is found with a *critical* violation, it fails the inspection. The agents’ objective is to inspect the establishments that have a high likelihood of failing an inspection early, thereby, preventing the establishment from serving harmful food for long.

Efficient Allocation: To schedule the inspections with high likelihoods of *critical violation*, the Chicago Department of Public Health (CDPH), along with the help of internal and external organizations, developed a machine learning (ML) model to predict a score for establishments around the city [Schenk Jr., 2015]. The ML model was trained on a dataset of more than 18,000 inspections recorded over four years by the department agents from Sep 2011 to Oct 2014 and made publicly available⁸. The dataset contained additional information like the weather on the inspection day, reports of crime and sanitation in the neighboring area. The department anonymized the identity of the agent who conducted each inspection in the dataset by clustering the agents into groups. The groups were created based on the agent *critical violation rate*. *Critical violation rate* can be computed as the ratio of the number of inspections with critical violations to the number of inspections completed. This signals how strict is an agent group. There were six of these groups, named after Chicago Transit Authority train lines – Purple, Blue, Orange, Green, Yellow, and Brown. Using all such features, the ML models output a likelihood score for the inspection. The

score is then used to order the inspections given a time period, which creates the inspection *schedule*. To evaluate the model, CDPH had a two-month test period where agents did the inspections as usual, and then the model predicted scores for these inspections without the ground truth. Compared to the random schedule by agents, the schedule obtained by ordering the inspections based on the model predictions reduced the average time to detect *critical violations* by 7 days in the simulation study. We use this model as the *efficient allocation* in our setting. More details about the dataset and the experimental setup are described in the work of Schenk Jr. [2015]⁹.

Fair Allocation: Singh *et al.* [2022] studied the impact of the efficient allocation with geographic fairness across nine regions of the city in mind. They found that the proposed ML model learns the signal from the agent group as a proxy for the likelihood of finding a *critical violation* and orders the inspections conducted by the most strict group first while delaying the others. Since there is a correlation between the agent groups and geographic regions, the benefit of scheduling some inspections early at the cost of others puts some regions at a disadvantage. They propose several methods to create a fair schedule by treating the agent groups as the *protected group* for fairness-aware ML models. They also propose two methods that modify the way the predicted score is used to create the schedule while achieving Pareto-dominance on the efficiency-fairness tradeoff. We adapt their ‘Sanitarian-blind’ approach as our *fair allocation*. The approaches are described in detail in Singh *et al.* [2022]¹⁰.

Experimental Setup: Given the original dataset, we want to explore the effect of change in the number of agents on the efficiency and fairness of the proposed solutions. Singh *et al.* describe the inspections as being conducted by “around three dozen” agents. Using this, we calculate the number of inspections completed per agent per day based on the assumption of 35 agents. Then, for any number of agents, we can compute the number of days required to complete a schedule. We vary the number of agents between $\{n \in \mathbb{Z} \mid 2 \leq n \leq 35\}$ to get a wide range for observing the effects on objectives. We assume that each agent roughly conducts an equal number of inspections, and we use a uniform distribution to assign inspections to agents. We further assume that the number of daily inspections stays uniform. Under this assumption, we use stochastic rounding to split inspections into days. We use this process for all numbers of agents and simulate an inspection schedule. Since the dataset contains a fixed number of inspections to be completed, the change in the number of agents changes the schedule length – fewer agents would take longer to complete the same amount of inspections and vice versa. For evaluation, we use 16 two-month test windows (following Singh *et al.* [2022]) for cross-validation.

We use outcome-based group utilities and costs. For costs, we follow Schenk Jr. and define the cost of the group as the average number of days that elapse before restaurants with critical violations in the group’s region of the city are inspected. Thus, if all were inspected on the first day that

⁸Updated dataset: https://data.cityofchicago.org/Health-Human-Services/Food-Inspections/4ijn-s7e5/about_data. Fixed dataset used for the study: <https://github.com/Chicago/food-inspections-evaluation>

⁹GitHub: <https://github.com/Chicago/food-inspections-evaluation>

¹⁰GitHub: <https://github.com/shubhams/eaamo-fair-food>




Figure 8: The plots show the efficiency (left Y-axis) and the absolute fairness (right Y-axis) using the efficient allocation from Schenk Jr. [2015]. Both measures increase with the increase in agents (resources). Error bars represent the standard errors.




Figure 9: The plot showing efficiency and fairness, similar to Fig. 8 but uses relative fairness on the right Y-axis. Higher values on both Y-axes are better. Error bars represent standard error.

cost would be zero while if all were inspected on the last day of the second month it would be approximately sixty. With fewer than 35 agents it would take more than two months to complete all inspections, making the maximum possible cost correspondingly larger. For utility, we instead compute the number of days in the inspection window after a restaurant with a critical violation is inspected. This is equivalent to subtracting the cost from the maximum window length for any amount of agents. Thus, if all were inspected on the first day, that utility would be approximately the time it takes 2 agents to complete the schedule (~ 1200 days), while if all were inspected on the last day with two agents it would be zero.

Verifying Objectives: Keeping the allocation policy fixed to the *efficient allocation* [Schenk Jr., 2015], we plot the efficiency and absolute fairness in Fig. 8. The figure shows the effect of change in the number of agents (resources) on the efficiency and absolute fairness separately. At lower levels of resources (between 2-10 agents), both efficiency and fairness show a rapid increase, and slow down approaching the high resources levels (30-35 agents). This shows that the group utilities resemble the outcome-based utilities (Eq. 4). We show the same allocation by using relative fairness in Fig. 9. Since relative fairness is measured on $[0, 1]$ interval, it appears that efficient allocation achieves higher fairness at lower resource levels and stays comparatively the same as the resources increase. This confirms the objective functions we are using in this setting are good approximations of the efficiency (Eq. 2) and fairness functions (absolute and relative, Eqs. 3, 5) used in the stylized examples.