We thank all the reviewers for their thoughtful comments, they are much appreciated. We will first address some general comments and then respond to comments made by the individual reviewers.

[General]. Data and code. ([R1, R3, R4]) To facilitate reproducibility of our results we will release all our code (incl. synthetic data generation) upon acceptance. As the data includes confidential medical information, we can only share our data upon request. However, we will append a table with some descriptive statistics of our data (amount of patients/organs, mean and std of patient and organ covariates, disease cohort sizes, etc.), as well as an overview of the features used for patients and organs. **Evaluation using ITE.** ([R2, R3]) Should OrganITE be the only policy sharing an ITE model with evaluation, we agree with the reviewers that OrganITE would be at an unfair advantage. However, this is not the case: every (estimation-based) policy shares the same ITE estimator, this way the policies are evaluated solely based on the policy, not on how well survival is estimated (see lines 283-286). **Morality and ethics.** ([R1, R2]) We agree that this work has significant moral implications, though we want to emphasize that such implications also hold for other algorithms and even clinicians themselves. However, as was indicated by [R2], we will move some of our discussion from the supplemental material to the main text, as we are allowed a 9$^{\text{th}}$ page for a camera-ready version. Furthermore, as was suggested by [R1], we will acknowledge (in the introduction) that our method is no exception to ethical concerns.

[R1]. Organ assignment policy taking into account rarity of organ. We agree with the reviewer that our statement in the conclusion about our method being the first to take into account the rarity of an organ may have been too general. What we meant to emphasize is that our organ-assignment policy learns in a data driven way what the optimal organ is for each patient, i.e the organ yielding the highest ITE for the patient (by considering the high dimensionality of the organ features) as well as the rarity of this optimal organ, by calculating its probability using an estimated density of all organs. This probability is estimated from the data. Note that this approach is different from the matching based method proposed by Dickerson et al. (2014) which considers a rule-based approach for defining organ compatibility in terms of blood groups which subsequently determines the rarity of certain organs (i.e. $\mathbb{I}_{\{v_s(\mathbf{X})>\tau\}}$, where $v_s, \tau \in [0, 100]$ are a patient's sensitization level and threshold respectively, while we define rarity as a distribution, $p(\mathbf{O}^*_{\mathbf{X}})$). Similarly, the OPTN policies used in the US also use a rule-based approach to define the organ rarity in terms of its compatibility. We will revise our claim in the paper, include a discussion and cite the relevant works pointed out by the reviewer.

[R2]. Density estimation. We agree that using a VAE is more standard in literature. As such, we have rerun our experiments using a VAE and will include these results in the main text. From Figure 1 we see that OrganITE with a VAE compares favorably to OrganITE using a KDE. **Optimisation target and reinforcement learning (RL).** We acknowledge that the optimisation target is heuristic. While one (with some effort) could indeed use RL to optimise for total life-years directly, we are convinced our method offers some advantages over RL: (i) explicitly taking into account an ITE estimator (and other components) allows us to interpret decisions made by OrganITE, as well as ease debugging of peculiar suggestions; (ii) for RL we would consider choosing a patient as the action, and the organ as the state resulting in a very sparse action space (as the actions are constrained to the patients currently in $\mathcal{X}^t_Q$) with minimal control over state-transitions, all resulting in an extremely hard to learn policy from logged data. However, we wish to stress that our arguments do not make it *impossible* to use RL in this setting, though it would result in a different paper entirely. **Error bars.** Not including confidence intervals (CIs) in the tables was an oversight for which we apologize. We will augment all our results with a 95% CI, as we have in our final experiment (Figure 4).

[R3]. We believe all concerns are responded to in the [General] section above.

[R4]. Notation. On line 121: $\mathcal{X} \subset \mathbb{R}^d$ is the *set* of possible patients (a subset of the real-vector space); on line 122: $\mathbf{X} \in \mathcal{X}$ is a random variable representing *one* patient as a d-dimensional real vector (with $\mathbf{x}$ a realisation of $\mathbf{X}$). Using this notation, Assum. 2 means that *all patient covariates (not set of patients)* $\mathbf{X}$ *that affect the treatment assignment,* $\mathbf{O}^{\pi_{obs}}_{\mathbf{X}}$*, and potential outcomes,* $Y^{\mathbf{o}}$*, are observed.* **Validity of assumptions.** While our assumptions are standard practice in ITE literature, we agree with [R4] that some assurance on whether or not these assumptions hold is warranted as they are *only* verifiable through domain knowledge. As such, we will note that: organ transplantation is a highly monitored setting where the many variables in our data are decided upon by (highly trained) clinicians. Note that we will append a description of our data as was mentioned above. **Potential outcomes (POs).** We will cite Pearl [1] as on p.166 he shows POs are shown to be equivalent to his framework, allowing the use of $do(\cdot)$. **Matching.** Our policy matches organs to patients without replacement; see lines 128-133. **Architectural details.** (32, 16, 16) indicates a 3-layered, fully-connected network, having widths 32, 16, and 16, respectively, which we will note in the supplemental material. In Figure 3 we concatinate $\mathbf{x} = (x_1, ..., x_d)^T$ and $\mathbf{o} = (o_1, ..., o_e)^T$ as $(x_1, ..., x_d, o_1, ..., o_e)^T$, which we will clarify.

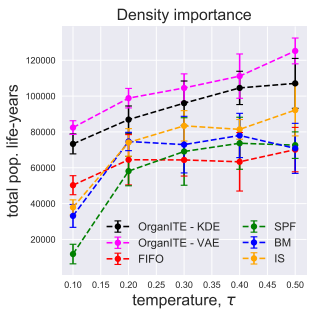

Figure 1: Density importance, with OrganITE - VAE

Reference: [1] Judea Pearl. Causal inference in statistics: An overview. *Statistics surveys*, 3:96–146, 2009.