

# *Ex-Ante Pareto and the Opaque-Identity Puzzle*<sup>\*</sup>

Johan E. Gustafsson and Kacper Kowalczyk<sup>†</sup>

ABSTRACT. Anna Mahtani describes a puzzle meant to show that the Ex-Ante Pareto Principle is incomplete as it stands and, since it cannot be completed in a satisfactory manner, decades of debate in welfare economics and ethics are undermined. In this paper, we provide a better solution to the puzzle which saves the Ex-Ante Pareto Principle from this challenge. We also explain how the plausibility of our solution is reinforced by its similarity to a standard solution to an analogous puzzle in quantified epistemic logic. We also show that even if the puzzle were to remain unsolved, its impact on welfare economics and ethics would be limited.

Two patients show up to your clinic: Ann, with arthritis, and Beth, with bronchitis. You have a single dose of a drug which can cure both ailments. While half a dose is better than nothing, it is much less effective than a full dose — so much so that it is better for a person to have a one-in-two chance of getting a full dose than to get a half one for sure. You can either give each patient half a dose or give a full dose to whoever arrived first. Since you think that Ann and Beth are equally likely to have arrived first, you reckon

(1) Ann's prospect is better if you give a full dose to whoever arrived first,

and

(2) Beth's prospect is better if you give a full dose to whoever arrived first.

You also know

(3) Ann and Beth are all the affected people.

<sup>\*</sup> Forthcoming in *The Journal of Philosophy*.

<sup>†</sup> We would be grateful for any thoughts or comments on this paper, which can be sent to [johan.eric.gustafsson@gmail.com](mailto:johan.eric.gustafsson@gmail.com).

From (1)–(3), you conclude

- (4) The prospect of everyone affected is better if you give a full dose to whoever arrived first.

Now, the receptionist tells you that Smith arrived before Jones. But, just as you are about to write a prescription for a full dose, you realize that you forgot Smith's and Jones's first names. You only remember that *either* Ann is called 'Smith' and Beth is called 'Jones' *or* it is the other way around. Since you know that Smith arrived first (and hence would get the full dose if you were to give a full dose to whoever arrived first), you reckon

- (5) Jones's prospect is not better if you give a full dose to whoever arrived first.

You also know

- (6) Jones is one of the affected people.

From (5) and (6), you conclude

- (7) It is not the case that the prospect of everyone affected is better if you give a full dose to whoever arrived first.

Since (7) contradicts (4), you have been led into a contradiction. This is *the Opaque-Identity Puzzle*.<sup>1</sup>

Anna Mahtani's resolution is that a claim like (1) is incomplete and must be understood as short for the following, partly metalinguistic, claim:

- (8) Ann's prospect (when designated as 'Ann') is better if you give a full dose to whoever arrived first.<sup>2</sup>

<sup>1</sup> Compare Anna Mahtani, "The *Ex Ante* Pareto Principle," *The Journal of Philosophy*, CXIV, 6 (June 2017): 303–23, at 310–11, and "Frege's Puzzle and the *Ex Ante* Pareto Principle," *Philosophical Studies*, CLXXVIII, 6 (June 2021): 2077–100, at pp. 2080–1. Similar examples were discussed in connection with an attempt combine non-utilitarianism with the idea of acting in everyone's *ex-ante* interest; see Elizabeth Harman, "Review of Caspar Hare, *The Limits of Kindness*," *Ethics*, CXXV, 3 (April 2015): 868–72, at p. 870, Caspar Hare, "Should We Wish Well to All?," *The Philosophical Review*, CXXV, 4 (October 2016): 451–72, at pp. 467–71, and Kieran Setiya, "Ignorance, Beneficence, and Rights," *Journal of Moral Philosophy*, xvii, 1 (February 2020): 56–74, at pp. 79–82, and "Other People," in Sarah Buss and Nandi Theunissen, eds., *Rethinking the Value of Humanity* (New York: Oxford University Press, 2023), pp. 314–36, at pp. 319–324.

<sup>2</sup> See Mahtani, "The *Ex Ante* Pareto Principle," at pp. 307–9, and "Frege's Puzzle," at pp. 2078–80.

Likewise, a claim like (4) is incomplete, since it does not specify how the value of the bound variable is to be designated and must be understood as short for one of the three possible versions of the following, partly metalinguistic, claim:

- (9) For every/some/special way of designating the people affected, the prospect of everyone affected (when so designated) is better if you give a full dose to whoever arrived first.

Mahtani concludes that any principle that involves quantified claims like (4) is likewise incomplete. This includes the *Ex-Ante* Pareto Principle:

*The Ex-Ante Pareto Principle* If the prospect of everyone affected is better if you choose one option than if you choose another, then you should not choose the second option if you can choose the first.<sup>3</sup>

Mahtani then argues that the *Ex-Ante* Pareto Principle is either implausibly strong or unhelpfully weak on any of her three readings of quantification.<sup>4</sup> This, she says, undermines decades of debate in welfare economics and ethics which have relied on the *Ex-Ante* Pareto Principle.<sup>5</sup>

We argue that there is a better solution to the Opaque-Identity Puzzle (§1). This solution, moreover, does not undermine the *Ex-Ante* Pareto Principle (§2). We show that the plausibility of our solution is reinforced by its similarity to a standard solution to an analogous puzzle in quantified epistemic logic (§3). Finally, we argue that the puzzle, even if it were

<sup>3</sup> A version of this principle was stated by John C. Harsanyi “Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility,” *Journal of Political Economy*, LXIII, 4 (August 1955): 309–21, at p. 313.

<sup>4</sup> See Mahtani, “The *Ex Ante* Pareto Principle,” at pp. 312–18; Mahtani, “Frege’s Puzzle,” at pp. 2089–94. Strictly speaking, Mahtani’s statement of the *Ex-Ante* Pareto Principle differs from ours in that it is ‘strong’ rather than ‘weak’, and in that it builds in the controversial Bernoulli Hypothesis, according to which a person’s prospect is better than another if and only if it offers that person a greater expectation of welfare; see John Broome, *Weighing Goods: Equality, Uncertainty and Time* (Oxford: Blackwell, 1991), p. 142.

<sup>5</sup> Mahtani, “Frege’s Puzzle,” at pp. 2096–9, also explains how her puzzle casts doubt on other related principles, such as Weak Pareto for Equal Risk, as in Marc Fleurbaey, “Assessing Risky Social Situations,” *Journal of Political Economy*, CXVIII, 4 (August 2010): 649–80, at p. 656, and principles favouring ex-ante equality, as in Peter A. Diamond “Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility: Comment,” *Journal of Political Economy*, LXXV, 5 (October 1967): 765–6.

left unsolved, would be limited in its impact on welfare economics and ethics (§4).

## 1. Our Solution

The probabilities in the Opaque-Identity Puzzle should be understood as subjective credences rather than as objective chances. There is no objective chance that your patients exchanged their identities. But you have some subjective credence that you mixed up them up.

Following Mahtani, we will assume that credences attach to epistemic possibilities, or epistemically possible worlds, whereas chances attach to metaphysical possibilities, or metaphysically possible worlds. Epistemic possibilities are ways the world might turn out to be for all we know, whereas metaphysical possibilities are ways the world could have been. For our purposes, the chief difference is that the former, but not the latter, can differ with respect to identity facts.<sup>6</sup>

For example, we now know that Hesperus is Phosphorus and, so, that Hesperus could not have failed to be Phosphorus. Thus there is no metaphysical possibility wherein Hesperus is distinct from Phosphorus. But, while ancient Greeks knew that Hesperus is Hesperus, they failed to realize that Hesperus is Phosphorus. Thus, for them, there was an epistemic possibility wherein Hesperus is Phosphorus and another epistemic possibility wherein Hesperus is distinct from Phosphorus.<sup>7</sup>

A solution to the Opaque-Identity Puzzle emerges once we realize that a person's prospect is like a probability-weighted portfolio of what can happen to that very person in different possible worlds. Accordingly, in order to determine a person's prospect, we need a method of trans-world identification.<sup>8</sup>

<sup>6</sup> This distinction is discussed in, among others, David J. Chalmers, "The Nature of Epistemic Space," in Andy Egan and Brian Weatherson, eds., *Epistemic Modality* (Oxford: Oxford University Press, 2011), at pp. 60–2. It appears that, without it, the puzzle cannot even arise. Indeed, a similar reaction to Mahtani's puzzle is discussed by Timothy Williamson, "Epistemological Consequences of Frege Puzzles," *Philosophical Topics*, XLIX, 2 (Fall 2021): 287–320, at pp. 311–14.

<sup>7</sup> The example derives from Gottlob Frege "Über Sinn und Bedeutung," *Zeitschrift für Philosophie und Philosophische Kritik*, c, 1 (1892): 25–50, at p. 32, translated by Max Black as "Sense and Reference," *The Philosophical Review*, LVII, 3 (May 1948): 209–30, at p. 215.

<sup>8</sup> More precisely, a person's prospect specifies what happens to that person in different possible states of the world, which can be equipped with probabilities. Broome also

In our initial example, there are two epistemically possible worlds, which can be qualitatively described as follows with two individuals in each world —  $i_1$  and  $i_2$  in the first world and  $i_3$  and  $i_4$  in the second, with no information about trans-world identities:

	first world		second world
$i_1$ :	called ‘Ann Smith’, first-comer, arthritic	$i_3$ :	called ‘Ann Jones’, runner-up, arthritic
$i_2$ :	called ‘Beth Jones’, runner-up, bronchitic	$i_4$ :	called ‘Beth Smith’, first-comer, bronchitic

We can distinguish two possible trans-world identifications. One is to identify  $i_1$  in one world with  $i_3$  in the other world and identify  $i_2$  in one world with  $i_4$  in the other world — that is, to identify by first name and medical condition. The other is to identify  $i_1$  in one world with  $i_4$  in the other world and identify  $i_2$  in one world with  $i_3$  in the other world — that is, to identify by last name and the time of arrival at the clinic.

Relative to the first identification, if you give a full dose to whoever arrived first, the people involved can have one of the following two prospects:

	first world (probability $\frac{1}{2}$ )	second world (probability $\frac{1}{2}$ )
first prospect	arriving first, getting a full dose	arriving second getting nothing
second prospect	arriving second, getting nothing	arriving first, getting a full dose

Either of these prospects is better than the prospect of half a dose for sure. So, relative to the first identification, the prospect of everyone affected is better if you give a full dose to whoever arrived first.

But, relative to the second identification, if you give a full dose to whoever arrived first, the people involved can have one of the following two prospects:

---

likens prospects to portfolios of outcomes; see Broome, *Weighing Lives* (Oxford: Oxford University Press), at p. 30. This is obscured in Mahtani’s discussion, as she uses ‘prospect’ in a non-standard way to mean what we would call ‘the expected welfare of a prospect’.

	first world (probability $\frac{1}{2}$ )	second world (probability $\frac{1}{2}$ )
first prospect	arriving first, getting a full dose	arriving first, getting a full dose
second prospect	arriving second, getting nothing	arriving second, getting nothing

Since half a dose is still better than nothing, it follows that, relative to the second identification, it is not the case that, if you give a full dose to whoever arrived first, the prospect of everyone affected is better.

Accordingly, whether the prospect of everyone affected is better depends on which method of trans-world identification is employed. So there is no need to relativize prospects to designators, as Mahtani does. They must already be understood as relative to trans-world identification methods. The appearance of contradiction arises only because no such method has been singled out in the example with which we started.<sup>9</sup>

## 2. The *Ex-Ante* Pareto Principle Rescued

Our solution, like Mahtani's own, has implications for how the *Ex-Ante* Pareto Principle should be understood and deployed. While the principle is never incomplete, the description of an example to which it is to be applied may be so. How is this missing information supplied?<sup>10</sup>

<sup>9</sup> Mahtani, "Frege's Puzzle," at p. 2087, considers a similar solution, according to which "we should focus on designators that pick out the same person at every state", which, she says, "may relate to the idea that some but not all designators are rigid". Her subsequent objection to that solution — that it is difficult to draw a principled distinction between rigid and non-rigid designators in the context of epistemic possibility — does not apply to us since our solution does not attempt to discriminate between different designators.

<sup>10</sup> Indeed, a puzzle analogous to the Opaque-Identity Puzzle arises as soon we introduce the idea of a person's prospect defined over epistemic possibilities, independently of issues raised by quantification over people's prospects and, so, independently of issues raised by the *Ex-Ante* Pareto Principle. In our example, it would be natural to believe

(10) Ann's prospect is better if you give a full dose to whoever arrived first,

but also

(11) Jones's prospect is not better if you give a full dose to whoever arrived first.

Then we seem to have

(12) Ann is not Jones.

No such method is implicit in the description of the example with which we started. It cannot be inferred, for instance, that Ann in one epistemic possibility is the same individual as Ann in the other epistemic possibility, that Smith in one epistemic possibility is the same individual as Smith in the other epistemic possibility, and so on. To see why not, recall that, in that example, there must be an epistemic possibility wherein Ann is the same individual as Smith and an epistemic possibility wherein Ann is one individual and Smith is another. It would then follow that the single individual appearing in the first epistemic possibility is identical with two distinct individuals appearing in the second epistemic possibility, which would contradict the transitivity of identity. Hence, given the symmetry of the case, it seems that neither Ann nor Smith in the second epistemic possibility can be identical with the individual in the first epistemic possibility. Thus no trans-world identities can plausibly be inferred on the basis of the names used in the description of the case.<sup>11</sup>

It seems that individuals populating one epistemic possibility will not, in general, also populate other epistemic possibilities. It is sometimes said that identity across epistemic possibilities should be understood in terms of a counterpart relation, not literal identity. An individual's counterpart in another world is that individual which resembles it closely enough in important respects and more closely than other individuals in its world.<sup>12</sup>

---

By an analogous argument,

(13) Beth is not Jones.

But the description of our example entails

(14) Jones is either Ann or Beth,

which contradicts the conjunction of (12) and (13). This puzzle is analogous to the veiled-man paradox often attributed to Eubulides of Miletus. See Lucian, *Vitarum Auctio* 22-3, translated by A. M. Harmon as "Philosophies for Sale," in *Lucian Volume II* (Cambridge, MA: Harvard University Press, 1915), pp. 451-511, at p. 495, Diogenes Laertius, 2.108, translated by Pamela Mensch as *Lives of the Eminent Philosophers* (New York: Oxford University Press, 2018), at p. 113, and Graham Priest, "The Hooded Man," *Journal of Philosophical Logic*, xxxviii, 5 (October 2002): 445-67.

<sup>11</sup> A similar argument can be found in Chalmers, "The Nature of Epistemic Space," at p. 84, but see also Mahtani, "Frege's Puzzle," at pp. 2087-8. Chalmers proposes an account of epistemic possibilities — which he calls 'scenarios' — according to which "[t]here are objects in scenarios, and they have properties, but only the properties are re-identifiable across scenarios (and here only some of them), and not the objects (except for abstract objects)"; see Chalmers, "The Nature of Epistemic Space," at p. 88.

<sup>12</sup> For example, after he argues that a wooden table could not have been made of ice, Kripke asks

Now, since resemblance is notoriously context-dependent, the operative counterpart relation will likewise vary between contexts. Thus, if we understood identity across epistemic possibilities in terms of a counterpart relation, it would follow that facts about which individuals face which prospects will depend on context. It would likewise follow that it will depend on context whether the antecedent of the *Ex-Ante* Pareto Principle holds or not. But, far from being a drawback, this sensitivity to context would allow us to better explain the pull of the Opaque-Identity Puzzle, in a way that is not otherwise ethically suspicious.

First, we could claim that there is a shift in the operative counterpart relation mid-puzzle. In the first half, the operative counterpart relation would give more weight to sharing a medical condition, while the counterpart relation operative in the second half would give more weight to arriving at the clinic at the same time. As a result, each claim involved in the puzzle would come out true relative to the counterpart relation operative in the context in which it initially appeared plausible.

Second, we could explain our inability to truly substitute ‘Jones’ for ‘Ann’ in claims about Ann’s prospects, even if Ann is in fact Jones. The occurrence of ‘Ann’ could well prompt a different context from that prompted by the occurrence of ‘Jones’. Then the property attributed to the single individual that is Ann Jones, say, would be different each time and, so, the truth-value of the attribution could likewise be different.<sup>13</sup>

Third, we could capture the popular intuition that extra information can affect which prospects should be considered in a given situation.<sup>14</sup>

---

What, then, does the intuition that the table might have turned out to have been made of ice [...] amount to? I think that it means simply that there might have been a *table* looking and feeling just like this one [...] which was in fact made of ice. [...] Something like counterpart theory is thus applicable to the situation...

See Saul Kripke, “Naming and Necessity,” in Donald Davidson and Gilbert Harman, eds., *Semantics of Natural Language* (Dordrecht: Reidel, 1972), pp. 253–355, at pp. 332–3. This stands in contrast with Kripke’s criticism of Lewis’s counterpart theory as applied to metaphysical modality; see Kripke, “Naming and Necessity,” at pp. 344–513, and David Lewis, “Counterpart Theory and Quantified Modal Logic,” *THE JOURNAL OF PHILOSOPHY*, LXV, 5 (March 1968): 113–26.

<sup>13</sup> Lewis noted a similar implication of his counterpart theory; see Lewis, “Counterparts of Persons and Their Bodies,” *The Journal of Philosophy*, LXVIII 7 (April 1971): 203–11, at p. 211.

<sup>14</sup> For this intuition, see Hare, “Should We Wish Well to All?” at pp. 466–71, and Setiya, “Ignorance, Beneficence, and Rights,” at pp. 69–73; “Other People,” at pp. 319–24. Mahtani, “Frege’s Puzzle,” at p. 2086, considers but in the end rejects the intuition at



Imagine, for example, that the receptionist at your clinic happens to locate thick dossiers on ‘Smith’ and ‘Jones’ but containing no first names. You find out a lot about the personal lives of Smith and Jones, for example, their education, projects, and careers. Even if beforehand it was natural to think of the situation in terms of what happens to Ann and Beth, it now seems natural to think of it in terms of what happens to Smith and Jones. This could again be explained by means of a shift in the operative counterpart relation, insofar as sharing the properties described in the dossiers is now more important for counterparthood.

But now imagine instead that you receive a phone call from Jones who apologizes for running a bit late before abruptly hanging up, without giving you a chance to ask for their first name. It appears that you find out nothing about this case that you did not know before. Now, it seems natural to think that someone’s prospect would be made worse by giving a full dose of the drug to whoever arrived first, so it now seems natural to think of the situation in terms of what happens to Smith and Jones. This would mean that a shift in the operative counterpart relation can also be prompted by a change in what is salient in a given situation, without any change in the information available.<sup>15</sup>

So, in this way, context could affect whether the antecedent of the *Ex-Ante* Pareto Principle holds. But none of this means that context could affect whether the consequent holds. Hence the context-sensitivity implied by the current proposal does not entail any troublesome context-sensitivity of claims about permissibility. To see this, note that the *Ex-Ante* Pareto Principle — always applied with respect to a trans-world identification supplied by context — is compatible with expectational utilitarianism, according to which the only permissible options are those which produce the greatest expected total of welfare. For example, consider that the verdict of expectational utilitarianism in our initial example remains unchanged, regardless of which trans-world identification is employed.

Indeed, if we accept the *Ex-Ante* Pareto Principle together with the hypothesis that the value of a person’s prospect is given by its expected welfare, then accepting expectational utilitarianism seems to be the only way to avoid context-sensitivity of claims about permissibility. So, if context-

---

issue. We can explain it without Hare’s appeal to incommensurability or Setiya’s appeal to the phenomenology of the face.

<sup>15</sup> We thank an anonymous referee for this observation.

sensitivity of permissibility claims is ruled out, our solution to the Opaque-Identity Puzzle could be used to streamline a version of Harsanyi's argument for utilitarianism.<sup>16</sup> On the other hand, if context-sensitivity is allowed, our solution implies that permissibility claims can be sensitive to changes in information or even — somewhat surprisingly — to changes in what is salient in a given situation.

Some dependence on context is unavoidable when applying the *Ex-Ante* Pareto Principle to prospects understood in terms of epistemic possibilities. This is because the context must provide a probability distribution across these possibilities. Moreover, if we are to evaluate epistemic modals more generally, context has to supply trans-world identity information, as we will see in the next section. Even if this information were missing in some situations, the *Ex-Ante* Pareto Principle would remain useful for establishing general theoretical results, such as those inspired by Harsanyi. Typically, a few suitably described hypothetical scenarios suffice for these purposes, with the necessary information being supplied by stipulation.

### 3. An Older Puzzle

A further advantage of our solution is that it is analogous to a standard solution to a similar puzzle about quantification into the scope of epistemic modals. The puzzle begins with the following example.<sup>17</sup>

You are investigating a murder. (You work for the police now.) And you know that the butler did it and that the gardener had nothing to do with it. So you reckon

(1\*) It is not the case that the gardener might have done it.

<sup>16</sup> See Harsanyi, "Cardinal Welfare," at pp. 312–4. We owe this point to Jake Nebel. Mahtani, "Frege's Puzzle," at pp. 2089–92, notes a similar but different connection between her discussion and Harsanyi's argument, as does Teruji Thomas, "The Asymmetry, Uncertainty, and the Long Term," *Philosophy and Phenomenological Research*, CVII, 2 (September 2023): 470–500, at pp. 480–1.

<sup>17</sup> Jelle Gerbrandy, "Identity in Epistemic Semantics," in Lawrence Cavedon, Patrick Blackburn, Nick Braisby, and Atsushi Shimojima, eds., *Logic, Language and Computation: Volume 3* (Stanford, CA: CSLI, 2000), pp. 147–59, at p. 147. The same example is discussed (albeit without quantification) in Gerbrandy, "Questions of Identity," in N. S. Care and R. M. Grimm, eds., *Proceedings of the Eleventh Amsterdam Colloquium* (Amsterdam: University of Amsterdam, 1997), pp. 133–8, at pp. 137–8.

The butler and the gardener are alone in a small room at the police station.  
So you conclude

(2\*) Not everyone in the room might have done it.

As you are about to write up a police report, you realize that you forgot the butler's and the gardener's names. You know that one is called 'Ann', the other 'Beth', but you forgot who is called what. So, for all you know,

(3\*) Ann might have done it,

and

(4\*) Beth might have done it.

Moreover, you know

(5\*) Ann and Beth are all the people in the room.

From (3\*)–(5\*), you conclude

(6\*) Anyone in the room might have done it.

Since (6\*) contradicts (2\*), you have been led, by plausible steps, into a contradiction. This is puzzling.

It should be clear that this puzzle is analogous to the Opaque-Identity Puzzle.<sup>18</sup> A standard solution to the former involves the idea of context-sensitive methods of trans-world identification. It is analogous to our own solution to the latter. For example, Gerbrandy gives a counterpart-theoretic semantics for the epistemic modal 'might'.<sup>19</sup> Aloni, on the other hand, keeps the standard semantics fixed and takes quantifiers to range over a contextually-specified subset of so-called individual concepts, so that trans-world identification methods are, in effect, built into the domain of the quantifiers.<sup>20</sup>

<sup>18</sup> The analogy was pointed out in Kacper Kowalczyk, "People in Suitcases," *Journal of Moral Philosophy*, xx, 1-2 (March 2023): 3–30, at p. 6n11.

<sup>19</sup> See Gerbrandy, "Identity in Epistemic Semantics." A similar account was developed by Dilip Ninan, "Quantification and Epistemic Modality," *The Philosophical Review*, cxxvii, 4 (October 2018): 433–85, at pp. 452–6.

<sup>20</sup> See Maria Aloni, *Quantification under Conceptual Covers* (Amsterdam: IILC, 2001), at pp. 105–8. Aloni's account is inspired by Jaakko Hintikka, "On the Logic of Perception," in Norman S. Care and Robert M. Grimm, eds., *Perception and Personal Identity* (Cleveland, OH: The Press of Case Western Reserve University, 1969), pp. 140–73, at pp. 160–73, and "Semantics for Propositional Attitudes," in J. W. Davis, D. J. Hockney, and W. K. Wilson, eds., *Philosophical Logic* (Dordrecht: Reidel, 1969), pp. 21–45, at pp. 30–43.

Mahtani's solution to the Opaque-Identity Puzzle is analogous to a different, less plausible, solution to this other puzzle, according to which a claim like (3\*) is incomplete and must be understood as short for the following, partly metalinguistic, claim:

(7\*) Ann (when designated as 'Ann') might have done it.

This would mean, presumably, that 'Ann' picks out Ann and, in some possible world, it picks out someone who has done it. Likewise, a claim like (6\*) would be seen as incomplete, as it does not specify how the value of the bound variable is to be designated, and it must be understood as short for one of the three possible versions of the following, partly metalinguistic, claim:

(8\*) For every/some/special way of designating the people involved, everyone (when so designated) might have done it.

This would mean, presumably, that, for every/some/special way of designating the people involved, everyone is picked out by their assigned designator and, in some possible world, that designator picks out someone who has done it.<sup>21</sup>

One problem with this solution is that it threatens to trivialize quantified epistemic modal claims, much like Mahtani's own solution to the Opaque-Identity Puzzle threatened to trivialize quantified claims about people's prospects. This is because designators are easy to gerrymander. For example, in typical cases where a murder has been committed, any individual can be assigned a designator that, in other possible worlds, picks out someone who has done it or, alternatively, someone who has not done it. As a result, the 'every' version of (8\*) will typically be trivially false, while the 'some' version of (8\*) will typically be trivially true.<sup>22</sup> And the outlook for the search for everyone's special designator is as bleak here as it was in the case of the Opaque-Identity Puzzle.

<sup>21</sup> Kaplan once attributed to Quine a similar account of quantification into modals; see David Kaplan, "Transworld Heir Lines," in Michael J. Loux, ed., *The Possible and the Actual: Readings in the Metaphysics of Modality* (Ithaca, NY: Cornell University Press, 1979), pp. 88–109, at pp. 95–6.

<sup>22</sup> Mahtani, "Frege's Puzzle," at pp. 2084–9, provides similar examples. Kaplan, "Transworld Heir Lines," at p. 95, provides a general recipe for devising a designator that picks out a given individual in one world and any individual we please in other worlds: Pick  $a$  in possible world  $w$  and  $b$  in possible world  $v$ . Suppose that  $\phi$  is a sentence true in  $w$  but not in  $v$ . Now consider the designator 'the  $x$  such that  $((x = a) \& \phi) \vee ((x = b) \& \neg\phi)$ '. It picks out  $a$  in  $w$  and  $b$  in  $v$ . If desired, a proper name could also be introduced by

#### 4. Other Versions of the *Ex-Ante* Pareto Principle

The Opaque-Identity Puzzle does not even arise for many prominent versions of the *Ex-Ante* Pareto Principle. In Harsanyi's version, the *Ex-Ante* Pareto Principle would correspond to the principle that a distribution of objective chance is morally preferred to another if everyone affected prefers the former to the latter.<sup>23</sup> So Harsanyi's version differs from Mahtani's in two respects: first, it is about objective chance rather than subjective credence and, second, it is about preference rather than betterness more generally.

The Opaque-Identity Puzzle does not arise insofar as the first difference obtains. To see this, recall that we are following Mahtani in assuming that credences attach to epistemic possibilities whereas chances attach to metaphysical possibilities. Now, the puzzle involved two epistemic possibilities, one wherein Ann is the same individual as Smith and another wherein Ann and Smith are different individuals. As Mahtani herself concedes, this could not happen if these possibilities were metaphysical, rather than epistemic.

Likewise, the Opaque-Identity Puzzle does not arise insofar as the second difference obtains. To see this, consider a hypothetical preferential version of the Opaque-Identity Puzzle. Recall that, in our original example, it is better for a person to get a one-in-two chance of a full dose than to get half a dose for sure. Let us assume that everyone involved has a corresponding preference. And suppose, without loss of generality, that Ann is Smith and Beth is Jones. There are now four cases to consider, depending on the information possessed by the people involved.

If Beth knows that she is Beth and that she is Jones, then she will know that she is the runner-up. So she will want you to split the dose. Hence one false step in this alternative version of the puzzle would be the preferential version of (2), namely,

---

means of this description, as in Evans's example of 'Julius,' stipulated to refer to whoever invented the zip; see Gareth Evans, *The Varieties of Reference* (Oxford: Clarendon Press, 1982), at p. 31.

<sup>23</sup> See Harsanyi "Cardinal Welfare," at p. 313 — who, however, states the principle in terms of indifference rather than strict preference. For this interpretation of Harsanyi, see, among others, Philippe Mongin and Marcus Pivato, "Social Evaluation under Risk and Uncertainty," in Matthew D. Adler and Marc Fleurbaey, eds., *The Oxford Handbook of Well-Being and Public Policy* (New York: Oxford University Press, 2016), pp. 711–45, at p. 713.

(2\*\*) Beth prefers that you give a full dose to whoever arrived first.

If, on the other hand, Beth knows that she is Beth but does not know whether she is Smith or Jones, then, presumably, she will care about the fate of the person called ‘Beth’ rather than about the fate of the person called ‘Jones’.<sup>24</sup> Hence she will want you to give a full dose to whoever arrived first. Since, in this case, Beth is Jones, one false step in this alternative version of the puzzle would be the preferential version of (5), namely,

(5\*\*) Jones does not prefer that you give a full dose to whoever arrived first.

Likewise, if Beth knows that she is Jones but does not know whether she is Ann or Beth, then, presumably, she will care about the fate of the person called ‘Jones’ rather than about the fate of the person called ‘Beth’. Hence she will want you to split the dose. Since, in this case, Beth is Jones, one false step would again be (2\*\*).

Finally, if Beth has no idea whether she is Smith or Jones but also no idea whether she is Ann or Beth, then she will be as it were behind a veil of ignorance. Presumably, she will want you to give a full dose to whoever arrived first. Since, in this case, Beth is Jones, one false step would again be (5\*\*).

The situation is analogous in case Ann is Jones and Beth is Smith. Hence you can know, in advance, that the preferential version of the Opaque-Identity Puzzle has a false premise. This is not affected by your uncertainty since your uncertainty would not change which mental states are possessed by the people involved.<sup>25</sup>

<sup>24</sup> Lewis suggests the following kind of example: Lingens might wish that he finds a way out of a library but (not knowing he is Lingens) might not wish Lingens to find a way out of the library; see Lewis, “Attitudes De Dicto and De Se,” *The Philosophical Review*, LXXXVIII, 4 (October 1979): 513–43, at pp. 529–30.

<sup>25</sup> Mahtani, “The *Ex Ante* Pareto Principle,” at p. 308n7, suggests that a potential lesson of her original puzzle is that “we must focus exclusively on the prospects calculated by the people concerned (no matter how ill-informed about the likelihood of the different possible outcomes given various actions)”. A preferential version of the *Ex-Ante* Pareto Principle could require, however, that the preferences of the people involved are based on a shared probability function, informed by the best available evidence. This requirement is independently motivated as a response to impossibility theorems that arise in cases of divergent probability assessments; see Broome, *Weighing Goods*, at pp. 151–64.

Therefore, the Opaque-Identity Puzzle does not arise for Harsanyi's version of the *Ex-Ante* Pareto Principle. So it does not undermine Harsanyi's argument for utilitarianism, which relies on that version of the principle. As we saw earlier, the Opaque-Identity Puzzle does not undermine Mah-tani's own version of the *Ex-Ante* Pareto Principle either, since the idea of a person's prospect presupposes a method of trans-world identification which must somehow be supplied before the principle can be deployed.

We wish to thank Jake Nebel and Nikhil Venkatesh for valuable comments.