The Pragmatics of Multiparty Communication Gricean theories of pragmatics propose that listeners interpret utterances by reasoning about the communicative intentions of the speaker [3]. The speaker's communicative intentions may vary along a number of dimensions: the state of the world that they are trying to communicate, the speaker's degree of knowledgeability [5], and the QUD [4]. Different reasoning may also be required in multiparty communication [1], i.e. situations in which the speaker intends to communicate with multiple listeners simultaneously.

Focusing on a simple implicature, we present both experimental data and a formal model, to better understand the core dynamics of multiparty pragmatics. We argue that when a pair of listeners is addressed, each reasons not only about the speaker but about the other listener too, and that this can result in the defusal of scalar implicatures. More generally, our data suggest that understanding the pragmatics of an utterance addressed to multiple parties involves more than just understanding the pragmatics of each speaker-listener pair. We propose an extension of the Rational Speech Acts model (RSA), which has previously been used to formalize two-party inferences, to account for this reasoning.

A Naming Implicature Suppose that Margaret and Arthur both regularly eat together at a cafe called *Flour*. If Margaret says (1) to Arthur, it carries the implicature for Arthur that she is referring to a place other than *Flour*, the name of which, unlike *Flour*, is not in their common ground.

(1) I got you cookies from a place I like. (2) I got you cookies from *Flour*.

On a Gricean account, this implicature stems from the fact that if the cafe *were Flour*, Margaret would have said (2) or similar, on the assumption that she aims to be maximally informative.

Multiparty Hypothesis In natural language, an utterance like (1) may be addressed to multiple people, e.g. the speaker Margaret may address two listeners, Arthur and Diane. If the name *Flour* is unknown to Diane (and Margaret and Arthur know this), we hypothesize that Arthur will not draw (or draw to a lesser degree) the implicature that the cookies did not come from *Flour*. This is because Arthur can *explain away* Margaret's avoidance of naming *Flour*: Margaret will not use this name if she believes that one of her listeners (Diane) may interpret it incorrectly. This requires Arthur to reason not only about Margaret's beliefs about himself, but about Margaret's beliefs about Diane.

Experimental Design To test the hypothesis that Arthur will not draw the "not *Flour*" implicature when he knows that (1) is addressed to both him and Diane, we conduct the following experiment. Participants are shown a short comic strip (see figure (1)), in which Margaret mentions her interest in *Flour* to Arthur. Subsequently, Margaret is seen meeting Arthur (condition *1*), Diane (condition *2*) or both (condition *3*) and telling her interlocutor(s) that she has brought them cookies from a place she likes. We then ask participants whether they think the cookies are from *Flour* (*yes/no* forced choice). We gather 268 participants on Mechanical Turk and exclude 52 of these, who either fail a comprehension question or don't complete the experiment. Each sees a single trial. We preregistered at `http://aspredicted.org/blind.php?x=3fp2jy`.

Analysis of Results Results are shown in figure (2). Significance was analyzed using Fisher's exact test. In condition (1), participants drew the implicature (i.e. answered *no*) at a significantly higher rate (42%) than in condition *2* (8%; $p < 0.001$), or condition *3* (13%; $p < 0.001$). The first comparison demonstrates the baseline two-party implicature, while the second shows the defusal of the implicature when a second listener is added, supporting our hypothesis.

Modeling the Implicature Recent formal work in pragmatics (RSA) derives implicatures by modeling reasoning between idealized rational speakers and listeners [2]. Within this paradigm, we first model the two-party case, by introducing assignment functions, and then extend to the multiparty setting.

Two Party Case Our aim is to define a model of the listener Arthur, which captures his pragmatic inferences after hearing an utterance like (1). Arthur is modeled as a conditional probability distribution L_1^A (for *Arthur*) which receives an utterance $u \in U$, and returns a distribution over states of the world $w \in W$. Like the L_1 model proposed by [2] for scalar implicature, L_1^A is *nested*, reasoning about a model S_1^A (of Margaret addressing Arthur) which reasons about L_0^A, a model of Arthur which interprets utterances literally. The key difference to the standard L_1 is the incorporation of an assignment function, to handle the semantics of proper names.

Minimally, the set of utterances U is $\{u_{Flour}, u_{Somewhere}\}$, corresponding roughly to (2) and (1) respectively, so that $u_{Somewhere}$ means that the speaker got cookies from somewhere (i.e. an existential claim). The set of worlds W is $\{w_{Flour}, w_{\neg Flour}\}$, representing the listener's belief as to whether the cookies are or are not from *Flour*. Assignment functions g map from proper names to entities. In this case, there is only one proper name to consider, u_{Flour}. For simplicity, we assume a domain with only two entities, an entity e_{Flour} and an entity e_{ALT} which represents all other cafes. There are thus two possible assignment functions (see figure 4), the correct one in which u_{Flour} maps to e_{Flour} and the incorrect one where it maps to e_{ALT}. For instance, $[\![u_{Flour}]\!]^{u_{Flour} \mapsto e_{ALT}}(w_{Flour}) = False$.

(3) $L_0^A(w|u, g_A) \propto [\![u]\!]^{g_A}(w) * P_A(w)$

(4) $V_A(w, u, g_A) = log(L_0^A(w|u, g_A))$

(5) $S_1^A(u|w, g_A) \propto exp(V_A(w, u, g_A))$

(6) $L_1^A(w|u) \propto \sum_{g_A \in G} S_1^A(u|w, g_A) * P_A(w) * P_{G_A}(g_A)$

Given an assignment function g_A, L_0^A hears an utterance and updates their beliefs to exclude incompatible worlds. For instance, $L_0^A(w_{\neg Flour}|u_{Flour}, g = u_{Flour} \mapsto e_{Flour}) = 0$. S_1^A, assuming an assignment g_A for Arthur, chooses an utterance with the goal of informing L_0^A. Importantly, if S_1^A assumes g_A maps the name u_{Flour} to the entity e_{Flour}, they prefer saying u_{Flour} to $u_{Somewhere}$ when they want to communicate that they got cookies from Flour, since it is more informative (for $g_A = (u_{Flour} \mapsto e_{Flour})$, $S_1^A(u_{Flour}|w_{Flour}, g_A) > S_1^A(u_{Somewhere}|w_{Flour}, g_A)$). Finally, L_1^A hears an utterance u and jointly infers what world S_1^A must have been in to have said u *and* what assignment function S_1^A assumed L_0^A was using.

P_{G_A} (for *Arthur*) represents L_1^A's prior beliefs about what assignment function S_1^A attributes to L_0^A. To represent the fact that the name of *Flour* is common ground between Margaret and Arthur, we set $P_{G_A}(g = u_{Flour} \mapsto e_{Flour}) = 1.0$ (since Arthur knows that Margaret knows that he knows the name). On hearing $u_{Somewhere}$, L_1^A prefers world $w_{\neg Flour}$, because S_1^A would have said u_{Flour} if *Flour* had been the cafe in question ($L_1^A(w_{\neg Flour}|u_{Somewhere}) = 0.75$). This corresponds to the calculation of the implicature from (1). (By contrast, an analogous model L_1^D of Diane, who does not know the meaning of u_{Flour} and therefore has a uniform probability over assignment functions, draws no implicature.) Throughout, we assume all priors over worlds and assignments are uniform but note that the qualitative predictions are robust to changes in prior beliefs.

Multiparty Model To extend the model to the multiparty setting, the key addition is a model of a speaker who wants to communicate to two listeners simultaneously. The change here is mathematically simple: when the speaker wants to communicate to multiple listeners, the utility of an utterance is the sum of utilities from communicating with each individual listener. For modeling condition *3* in the experiment, in which Margaret is addressing both *Arthur* and *Diane*, we define S_1^{AD} (Equation 11). In this definition, Margaret gains utility from communicating to both L_0^A (*Arthur*) and L_0^D (*Diane*).

(7) $L_0^D(w|u, g_D) \propto [\![u]\!]^{g_D}(w) * P_D(w)$

(8) $V_A(w, u, g_A) = log(L_0^A(w|u, g_A))$

(9) $V_D(w, u, g_D) = log(L_0^D(w|u, g_D))$

(10) $V_{AD}(w, u, g_A, g_D) = V_A(w, u, g_A) + V_D(w, u, g_D)$

(11) $S_1^{AD}(u|w, g_A, g_D) \propto exp(V_{AD}(w, u, g_A, g_D))$

(12) $L_1^{AD}(w|u) \propto \sum_{g_A, g_D \in G} S_1^{AD}(u|w, g_A, g_D) * P_A(w) * P_{G_A}(g_A) * P_{G_D}(g_D)$

S_1^{AD} assumes a fixed assignment function g_A for L_0^A and g_D for L_0^D. L_1^{AD}, reasoning about S_1^{AD}, now jointly infers *both* g_A and g_D, as well as the world w. As before, $P_{G_A}(g = u_{Flour} \mapsto e_{Flour}) = 1.0$ (*Flour* is common ground between Arthur and Margaret) but since Diane is not familiar with *Flour*, $P_{G_D}(g = u_{Flour} \mapsto e_{Flour}) = P_{G_D}(g = u_{Flour} \mapsto e_{ALT}) = 0.5$. This account assumes that a name has a particular referent for L_0^D, but what this referent is may be unknown by the pragmatic listener L_1^A.

The listener L_1^{AD} can explain away $S_1^A S$'s choice to say $u_{Somewhere}$ (see figure 3): if the speaker believes that u_{Flour} will be misinterpreted by listener *Diana* L_0^D, then she will choose the less informative but also less risky utterance $u_{Somewhere}$. Thus $L_1^{AD}(w_{\neg Flour}|some) < L_1^A(w_{\neg Flour}|some)$, corresponding to the experimentally observed difference between conditions *1* and *3*.

Conclusions This work presents empirical data and an accompanying formal model of a simple case of multiparty pragmatics. Using an implicature involving proper names, we model a multiparty communication where the presence of a second listener defuses an implicature the first would otherwise draw.

Figure 1: Relevant panels of the comic shown in the experiment, with condition (1) in blue and condition (3) in red. (Condition (2) resembles condition (1) but with Diane, not Arthur.) Participants are asked: Do you think Margaret got the cookies at *Flour*?

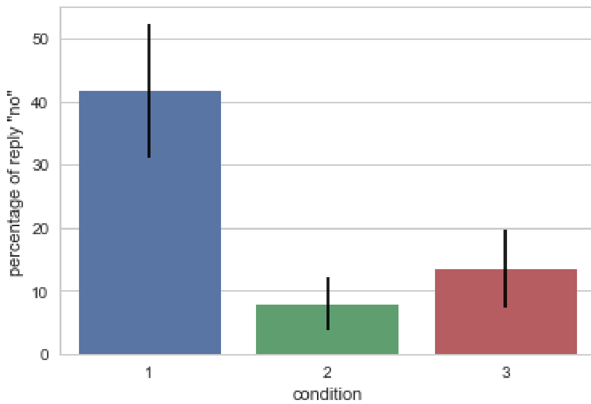

Figure 2: Experiment results for all conditions. Error bars are 95% confidence intervals.

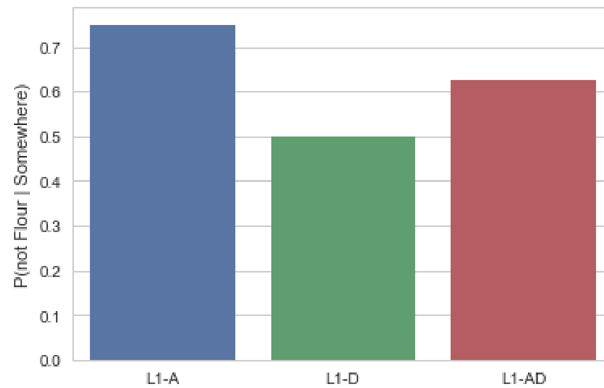

Figure 3: Model predictions for L_1^A, L_1^D, and L_1^{AD}. Qualitative predictions are robust to changes in model parameters: the implicature is strongest when the audience is only Arthur, weakest when it is only Diane, and intermediate when it is both.

w	$P_A(w)$	$P_D(w)$	g_D	$P_{G_D}(g_D)$	g_A	$P_{G_A}(g_A)$
w_{Flour}	0.5	0.5	$u_{Flour} \mapsto e_{Flour}$	0.5	$u_{Flour} \mapsto e_{Flour}$	1.0
$w_{\neg Flour}$	0.5	0.5	$u_{Flour} \mapsto e_{ALT}$	0.5		

Figure 4: Prior distributions over worlds and assignment functions. Arthur and Margaret have common knowledge that Arthur interprets *Flour* correctly. In contrast, there is uncertainty about whether Diane interprets the utterance correctly.

References

[1] Herbert H Clark and Thomas B Carlson. "Hearers and speech acts". In: *Language* (1982), pp. 332–373.

[2] Noah D Goodman and Michael C Frank. "Pragmatic language interpretation as probabilistic inference". In: *Trends in cognitive sciences* 20.11 (2016), pp. 818–829.

[3] H. Paul Grice. "Logic and Conversation". In: *Speech Acts*. Ed. by Peter Cole and Jerry L. Morgan. Vol. 3. Syntax and Semantics. New York: Academic Press, 1975, pp. 41–58.

[4] Craige Roberts. "Information structure in discourse: Towards an integrated formal theory of pragmatics". In: (1996).

[5] Uli Sauerland. "Scalar implicatures in complex sentences". In: *Linguistics and philosophy* 27.3 (2004), pp. 367–391.