UNIVERSITY OF CALIFORNIA, MERCED

From Microbes in Motion to Pumas in Pursuit of Prey: Foraging Phenomena in Interactive Multi-agent Systems

A dissertation submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy
in

Physics

by

VPS Ritwika

Committee in charge:

Dr. Bin Liu,Chair
Dr. Justin Yeakel
Prof. Ajay Gopinathan

2021

The dissertation of VPS Ritwika, titled From Microbes in Motion to Pumas in Pursuit of Prey: Foraging Phenomena in Interactive Multi-agent Systems, is approved, and it is acceptable in quality and form for publication.

(Prof. Ajay Gopinathan) Principal Adviser	Date
(Dr. Bin Liu) Committee Chair	Date
(Dr. Justin Yeakel) Committee Member	Date

University of California, Merced

2021

For me, because wow-I-actually-did-it, even through all the anxiety and depression and my brain screaming at me that I am a garbage grad student, and even though writing feels like pulling teeth.

Also for Mania, Bastet, Isis/Charlie, and Anubis, for all the ways you make my life better and richer.

Oh, and for anyone who needs to hear this: kids, do not repeat my mistakes. Start writing your dissertation in like, your second year.

Acknowledgments

There are so many people who have held my hand through the last several years and made this whole PhD thing possible, so I imagine this is going to be a few pages. In fact, I put this off for so long because I was terrified of forgetting someone—and I still have half a mind to just write *Thank you, you know who you are*, but I won't, because all of you deserve more. So here goes.

To start with, I want to say that I have known and/or heard about too many bright-eyed young scientists go through traumatising grad school experiences, and somehow, I had the incredible good fortune of having not one, but *two* wonderful advisors during my PhD. To top that, I somehow managed to also acquire several wonderful mentors and collaborators on the way. My gratitude, quite simply, is boundless.

Ajay and Justin, without an ounce of exaggeration, have been the best advisors I could have asked for. They helped me grow so much as a scientist and science communicator: I have gone from corrections on every slide in my first year presentations to being able to make slides with minimal edits. I have also gone from our one-on-one meetings being an hour of them guiding me every step of the way to me being able to figure things out on my own. And while I still need help and guidance with things, I am definitely not the floundering first year who didn't even know where to start with a project. That's all them.

But that's not all. They have been firmly in my corner all these years. They supported my wacky I-will-do-all-the-science-I-think-is-cool approach to research, helped me realise my vision for what I wanted out of grad school, provided mentorship tailored to my needs and anxieties, were incredibly kind and understanding whenever I needed a mental health break, fostered the kind of lab culture which has resulted in two wonderful research groups, put up with my unintentional tendency to get everything done at the last possible hour, and never made me feel stupid because I didn't know something. I would not be standing at the end of this road without them and I hope that they haven't gotten tired of me, because

I would LOVE to continue orbiting their scientific periphery (seriously, if there was a way for me to stay on with both of them, I would have. But seeing as how that isn't happening, I dearly hope that this is not the end of my time working with them!). Suffice it to say, I am so, so grateful for everything they have done and continue to do for me.

Anne Warlaumont, longtime collaborator and current (postdoc) advisor, has been a wonderful mentor, unwaveringly kind and patient, brilliant, relatable, and supportive. I want to thank Anne for being ok with my anxiety when it comes to handling sensitive data and with my need to be as thorough as possible when reporting results, for sharing (and encouraging) my outlook on open science and the value of null results, for her curiosity about my work outside of infant vocal behaviour and her support of my diverse research interests, for being understanding of my struggles with mental health and resultant delays with research, for making research meetings feel like catching up with a friend over coffee, for being supportive of me working from two time zones away, and in general, for leading by example.

Carrie Menke is such a pillar of UC Merced (and especially the Physics Department), and every semester I was a TA for Carrie has been so much fun. Carrie helped me grow as a teacher, listened to my ideas to improve students' learning experience, trusted me with students' education, listened to me whenever I came to her in need of advice, and has been one of my biggest cheerleaders at UC Merced. I am grateful to Carrie for all these things, and also for caring so deeply about her TAs and students, and their experience in the classroom. I am also grateful to Toni Stone for her kindness and support during the semester I was a TA for her. She made the experience a fun one and had my back when I was treated unfairly by others during my time as her TA.

Suzanne Sindi is yet another non-advisor mentor figure I am grateful to and one of those people who has made UC Merced so special. While I never worked with Suzanne in any official capacity, she has been so welcoming and open since the first day I met her in the COB hallway. I am thankful to Suzanne for adopting me into the COB Applied Math cohort over the course of many Applied Math coffee hours, for always having a kind word for me, for understanding and supporting my advocacy work at UC Merced, and for caring deeply about mental health and inclusivity in academia.

(And on that note, during my first few years at UC Merced, when my office was in COB, the Applied Math folks made up a big chunk of my professional social circle, and I am grateful to them for welcoming me with open arms.)

Bin Liu, who is the Chair of my committee, has been so very patient with me over the years, and especially these last few months, as I struggled with writing this dissertation and missed deadline after deadline (writing is, by a wide margin, the narrowest bottleneck I have encountered with research). And before that, he always had valuable suggestions with my work, and has been so very kind to me from the time he taught Electrodynamics to my cohort all the way till now.

Now, the fantastic thing about great advisors is that their groups tend to be made of wonderful people as well. And so, I have been fortunate enough to have not one, but *two*, great research groups I have had the privilege of being a part of, during my time at UC Merced.

The members of Ajay's research group—Katie, David, Farnaz, Imtiaz, Niranjan, Suraj, Jose and Chandra—have all played a part in making the Gopinathan group such a fun space, both research-wise and personally. I am especially grateful to David for freely sharing his seemingly boundless expertise with all things computational and network administrate-y; to Katie, for welcoming me into the group during my first year, when it was just me and her, for all the times she did the 'senior grad student mentoring junior grad student' thing so well, and for all her help whenever I got stuck with a particularly finicky piece of code; to Niranjan, for being such a joy to work with (we have an almost-finished set of computational modules to teach biophysics at the university level), and for his earnestness, sincerity, and friendship; and to Suraj, Jose, and Farnaz, for being such lovely friends.

Not to be outdone, the members of Justin's research group—JP, Uttam, Taran, Irina, and Megha—provided me with another wonderful and fun research space to exist in. All of them have made the transition from people I work(ed) with to cherished friends; I think part of it is the open ended format of our group meetings. There was this guaranteed weekly space for us to be ourselves in a group and connect, which is another thing I am grateful to Justin for. I am thankful for their presence in my life and for all the ways they have helped me through grad school, be it freely sharing their research expertise or just listening to me when I sorely needed someone to do just that.

I have also been fortunate to have an inordinately large number of wonderful collaborators for the career stage I am at, which is yet another thing I attribute to Ajay and Justin (I hope you see the pattern here, of how so many of the good things about my grad school experience are attributable to them).

I had the good fortune of being a part of the 2019 SFI Complex Systems Summer School

(CSSS) and the 2019 Complex Networks Winter Workshop in Quèbec. They were the best, most concentrated, and most fun scientific experiences of my life, and I met some wonderful people there, a lot of whom ended up being collaborators on a number of fun projects. I am incredibly grateful for these experiences (thank you Justin, for convincing me to apply to CSSS) and in particular, would like to thank Henri Kauhanen, Jessica Audrey Lee, Anshuman Swain, Brennan Klein, Dakota Murray, Louisa Di Felice, Kate Wootton, Billy Braasch, Dan Larremore, Dave Feldman, JP Gonzales, and Phil Chodrow.

Chapter 2 in this dissertation is from a collaboration with KC Huang at Stanford. I have worked with KC for several years now, and I am grateful to KC for his kindness, grace, good advice, and good humour through the years, and for his incisive insights into the project that can only come from the experience of working with experiments. I am also grateful to Tristan Ursell and Rosanna Chau, who are KC lab alumni and collaborators on this project, for their generosity and kindness in sharing relevant code and answering questions when I had them.

Chapter 3 is the product of a collaboration with several wonderful people—Chris Kello, Anne (who I have already mentioned), Sara Mendoza, Gina Pretzer, and Chris Shedd—without whom this work would not have been possible. In addition to providing thoughtful and insightful comments on this work and sharing his expertise on foraging processes in the context of cognition, Chris has also been wonderful to work with in his capacity as the Interim Graduate Dean while I was part of the Graduate Student Association (GSA).

Speaking of Graduate Deans, I am extremely grateful to Marjorie Zatz, who was the Graduate Dean during most of my time at UC Merced. I cannot thank her enough for fostering a culture where I felt like I could talk to her about graduate student issues and for making my time as co-President of the GSA such a wonderful experience.

My time with the GSA—first as co-President and then as the International Affairs Officer—is one of the fondest memories I have of my time at UC Merced, and I cannot thank the cohort of officers I worked with enough for making it so. I am grateful to them for their unwavering dedication to the cause of graduate student welfare, their courage and conviction in the face of injustices, and their justified anger at things that should be better. In particular, I am grateful to Brandon for nominating me for the role of GSA President in the first place, and then, for being such a wonderful co-President during my first year with the GSA; and to Shayna, Jordan, and Taylor for being wonderful friends who consistently provided valuable feedback about my work with the GSA and for always going to bat for

graduate student welfare. And Taylor, thank you also for bringing Gracie into my life!

I am grateful to the leadership of the Merced Indian Graduate Student Association (MIGSA) and UC Merced's W-STEM chapter, for making the time I spent as an officer for both organisations a productive and wonderful experience. Together, my experience with the GSA, MIGSA, and W-STEM has been formative to my conviction that I want leadership and advocacy efforts to be an integral part of any career path I embark upon.

I would not have gotten through this PhD without the support, administrative, and health services staff at UC Merced. The physics lab support staff—Stefanie Stepp in particular—have consistently done a fantastic job making sure that the physics labs were fun to teach and addressed issues as soon as they were brought up. I am thankful for their dedication and rigour in making sure that the labs ran smoothly, week after week.

The folks at Grad Div and the Office of International Affairs—Cassie Gunter, Jen Quiralte, Lacey Long Vejar, Becky Mirza, Paul Roberts, Rita Guel, and Jesse Batther—have been incredibly helpful throughout my time as a grad student. I am thankful to them for being a reliable resource I could turn to when I needed help with the many documentation requirements as an international student and for answering any and all questions I had about the administrative side of getting a PhD. Thank you also to Carrie Kouadio, executive director of CCBM, for having my back every time I did CCBM virtual workshops with Ajay, and to Belinda Braunstein, for being such a joyful presence at UC Merced.

Linda Hirst, Sayantani Ghosh, and Chih-Chun Chien (they are definitely not administrative staff, but did quite a bit of administrative work in their capacity as Physics grad group Chairs, so this is where I am plopping this paragraph) have been fantastic Chairs of the Physics graduate group during my time as a graduate student, and I am grateful to them for supporting me (and other graduate students) during their tenure.

I am also incredibly grateful to the folks at the Student Health Center and CAPS, and especially Heather Weigel, Dr. Laurel Ealy, Amanda Peterson, Dr. Michelle Brinkhop, Alyson Holman, Martha Durand, and Michelle Law. They took such good care of my health and being able to not worry about healthcare in a country whose healthcare system is so convoluted was such a boon.

I am thankful to Marie Claire, for her boundless kindness while I struggled with my anxiety, and for telling me that it will all be OK, and for being so kind and generous with her advice and support (and application materials) while I was applying for postdoctoral fellowships.

I am also thankful for my time at Crossfit Merced for bringing Sid, Lindsey, Holly, Matt, and Gus into my life, who have taught me so much about fitness, educated me in the fundamentals of Olympic and power lifting (which has helped my mental health so much), and are wonderful people. I am especially grateful to Sid (and by extension, Bleu), for being such a lovely presence in my life and for being such a ferocious source of inspiration.

For the first four years in the US, I lived in a huge house with some of the best housemates I could ask for, and I can safely say that they have ruined me for all other housemates. Arrash, Eric, Sarang, Nivin, and Rachel made Jenner house such a special place, and I truly do not know if I would have survived my PhD had I ended up anywhere else during those first few years. Arrash and Eric, in particular, have become some of my closest friends from my time in grad school and I am grateful for your (and by extension, your families') presence in my life.

And now, in no particular order, is the enormous list of friends I have made during my time at UC Merced and to whom, I am endlessly grateful: Ekta and Anshika (for making the experience of living with the worst housemate I have ever had somewhat bearable); Mark, Lucy, and Ahmed (for being such fun, kind, and lovely friends); Amanda (for being arguably one of the nicest, kindest people I know, and for all the Christmas cookies); Majerle, Maia, and Shannon (for all the lovely Girls' Nights); Micah (for being such an incredible hype person and for always making me smile); Arabi, Abhinav, and Adityaa, (for being such an enormous help during my move from Merced and for generally being such wonderful friends); Palak (for being one of my closest friends in grad school, for all the walks and good food, and for meeting me at my level of weirdness); Noelle, Rob, and Xena (for being some of my sweetest friends, and especially Noelle, for validating and understanding my struggles with mental health); Ann (for everything you do for UC Merced grad students, and for being such a lovely, sweet friend); Akshat (for always having a place for me in the bay and for making me laugh with your truly horrendous puns); Judith (for keeping me accountable and for being one of my fiercest champions); Arya (for all those nights we have chatted away in Malayalam, for making me feel like I have a piece of home here, and for everything you mean to me); and John (for all the fun travels, for the kitties, and for the presence of your family in my life).

I have been fortunate enough to retain my closest friends all the way from middle school, despite the time zones separating us. They have cheered me on through undergrad and grad school, and I am so grateful to them. Arya, CD, Athira, and Krishna have been particularly

kind and understanding of my long absences from their lives and welcomed me back with open arms whenever I found the energy to go say hi again. I love you and I am so lucky to have you.

There is a wonderful cohort of people from my undergrad who moved to the US for grad school, and pre-pandemic, one of the things I looked forward to the most was opportunities to travel with or visit them. All of you kept me sane and I am grateful for your friendship. I am particularly grateful to Pranav (and Divya), for providing me with what I consider a home away from home in Chapel Hill, and for all the friends you helped me make there (Jonathan, Brenna, Ayushi, Sandeep, and Aatish, thank you for your friendship and for making every trip to Chapel Hill so much fun); to Animakshi (who is now in India), for always being everything I needed you to be; to Ram (alternatively, bestie) for the movie nights and Chicago trips and everything else; and to Pooja, Athira, Daniel, Tanvi, Saranya, and Dhanesh, for your kindness and friendship.

Looking back, I never would have gone to grad school without the support and mentorship of faculty from IISER-TVM, my undergraduate institution. I am grateful to Dr. Anil Shaji, Dr. Hema Somanathan, Dr. Joy Mitra, Prof. MS Gopinathan (who incidentally, is Ajay's father; funny how these things work), Dr. Deepshikha Nagar, and Prof. Subodh Shenoy, for their mentorship, advice, and support, without which I would not be here.

I am also incredibly grateful to my teachers from high school, who still ask me how my PhD is going and are generally so excited for all my milestones. In particular, I am grateful for the unwavering support from Remani teacher, Sudha teacher, Yamuna teacher, Anitha teacher, Shobha teacher, Usha teacher, and Shobhana teacher[1].

I also want to thank my dad and mom, my extended family and assorted family friends, who have been incredibly invested and supportive throughout my PhD. I am grateful for your patience, kindness, understanding, and support as I trudged my way through the past six years.

And finally, I am so, so grateful for Mania (I like to call her my soul dog; she belongs to Arrash, which means I have to thank him again here), and my kitties, Bastet, Charlie/Isis (thank you, Sofia, for letting us keep her after accidentally stealing/rescuing her from Merced; it is a long story, but suffice it to say that we have her blessing), and Anubis, for everything they do to add joy to my life.

[1]I have always addressed them this way, with their name followed by 'teacher' and it is not a habit I think I will ever break

I truly hope that I haven't forgotten anyone here. But if I have, please know that it is because my brain is extremely porous right now, and not because you don't matter. You do, so very much, and I am so thankful for your presence in my life.

VPS Ritwika

Education

2015–present **PhD candidate (Physics)**, *University of California, Merced*, GPA: 3.9/4.
Advanced to candidacy: Fall 2018 (10/2018); Graduation: expected June 2021

Advisor(s) **Prof. Ajay Gopinathan (Physics, UC Merced) and Dr. Justin Yeakel (Life and Environmental Sciences, UC Merced)**.

2010–2015 **Int. BS-MS (Physics)**, *Indian Institute of Science Education and Research, Thiruvananthapuram (IISER TVM), India*, *GPA: 8.89/10*.
Thesis: Environmental Effects in Excitonic Energy Transfer

Academic and Professional Experience

Graduate Student Researcher: Summer 2016, Spring 2017. Funded under grant titled 'Infant Vocalization as Foraging for Caregiver Responses'. Funding Agency: NSF; Fall 2017 (25 %), Spring 2018, Fall 2018 (25 %), Spring 2019 (25 %), Fall 2019 (25 %) , Summer 2020

Teaching Assistant: Fall 2015, Spring 2016, Fall 2016, Fall 2017 (25 %), Fall 2018 (25 %), Spring 2019 (25 %), Fall 2019 (25 %) , Spring 2020, Fall 2020

Course developer: Developed computational course modules for PHYS 204, Biophysics at UC Merced

CCBM Virtual Sessions, Summer/Fall 2020: Facilitated workshops titled 'Understanding Living Systems through Scientific Computing' and 'The Science of Flocks and Swarms' as part of the Center for Cellular and Biomolecular Machines Summer/Fall 2020 Virtual Sessions, with Prof. Ajay Gopinathan

Awards and Fellowships

2021 **Steve and Mia Kang Student Leadership for Social Change Prize**, *Margo F. Souza Student Leadership Center, UC Merced*.

2021 **Building Community Award**, *Margo F. Souza Student Leadership Center, UC Merced*, (awarded to the UC Merced Graduate Student Association).

2021 **Outstanding Student Service and Outreach Award**, *Department of Physics, UC Merced*.

2021 **UC Merced GradSLAM Top 10 Finalist**, *UC Merced*.

2021 **Graduate Dean's Dissertation Fellowship**, *UC Merced*.

2020–2021 **CCBM Scholar**, *Center for Cellular and Biomolecular Machines (CCBM), UC Merced*.

2019, 2020 **GRAD-EXCEL Peer Mentorship Program Award**, *UC Merced*.

2019 **Graduate Fellowship Incentive Program Award**, *UC Merced*.

2019 **Travel award through the Dean's office**, *UC Merced*.

2018, 2019 **APS DBIO Student Travel Award**, American Physical Society (APS).

2018 **Broadening Participation Committee Travel Award**, Society for Integrative and Comparative Biology (SICB).

2018 **Charlotte Mangum Student Support Program Award**, SICB.

2017, 2018, 2019 **UC Merced Physics Summer Fellowship**, UC Merced.

2017, 2018, 2020 **UC Merced Physics Travel Award**, UC Merced.

2016 **Top performer in Physics prelims and Physics core courses**, UC Merced.

2015–2016 **Graduate Group Recruitment Fellowship**, UC Merced.

2010–2015 **INSPIRE Fellowship**, Department of Science and Technology, Govt. of India.

2013 **DAAD WISE Fellowship 2013**, DAAD (was unable to pursue).

Publications

1. **VPS Ritwika**, Gina Pretzer, Sara Mendoza, Tim Shea, Christopher Shedd, Christopher Kello, Ajay Gopinathan, Anne Warlaumont. **Exploratory dynamics of vocal foraging during infant-caregiver communication.** [1] Scientific Reports. 10(1):1–14, 2020.

Manuscripts in Preparation

* Indicates equal contribution

1. **VPS Ritwika**, Rosanna Chau, Devaki Bhaya, KC Huang, Ajay Gopinathan. **Initiation of fingering instabilities during cyanobacterial phototaxis.**
2. **VPS Ritwika**, Ajay Gopinathan, Justin D. Yeakel. **Fitness Trade-offs of Predation: When to Scavenge and When to Steal.**
3. **VPS Ritwika**, Sara Schneider, Lukas Lopez, Sara Mendoza, Christopher Kello, Ajay Gopinathan, Anne Warlaumont. **Infant Vocalization as a Foraging Process: Validating LENA labels with human labels.**

Conferences

03/2021	APS March Meeting (Talk)
03/2020	APS March Meeting (Talk, presented virtually in Physics of Social Interactions)
03/2019	APS March Meeting, Boston, MA
08/2018	Ecological Society of America (ESA) Annual Meeting, New Orleans, LA (Talk)
04/2018	Central Valley 2018 SIAM Regional Student Conference, UC Merced (Poster)
03/2018	APS March Meeting, Los Angeles, CA (Talk)
01/2018	Gordon Research Conference on Predator-Prey Interactions, Ventura, CA (Poster)
01/2018	SICB Annual Meeting, San Francisco, CA (Talk; also substitue co-chair of session Social Behavior)
11/2017	APS Far West Section Annual Meeting, UC Merced (Talk)
10/2017	UC Merced SIAM Student Chapter Seminar Series (Talk)
03/2017	APS March Meeting, New Orleans, LA (Talk)
02/2015	Young Quantum (YouQu), Harish-Chandra Research Institute (HRI), Allahabad, India (Poster)
12/2011	Young Ecologists Talk and Interact (YETI), IIT Guwahati, Assam, India (Poster)

Workshops and Schools

Spring 2021	Spring 2021 Institute for Evidence-Based Teaching Practices, UC Merced
03/2021	Mathematical and Computational Approaches to Social Justice (held virtually), The Institute for Computational and Experimental Research in Mathematics, Brown University
12/2019	Complex Networks Winter Workshop, Quebec City, Canada
11/2019	CCBM Workshop on Emerging Themes in Cellular and Biomolecular Machines, UC Merced
Summer 2017	CCBM-CREST Summer Workshop: Introduction to Scientific Computing
Summer 2019	Santa Fe Institue Complex Systems Summer School

Technical and Computer Skills

MATLAB, R, Mathematica, LaTeX, NetLogo, Python, Julia, Perl

Outreach and Service

05/20–present	International Affairs Officer, UC Merced Graduate Student Association
05/20–present	Graduate Treasurer, W-STEM, UC Merced
01/20–05/20	Graduate student member, Graduate Council, UC Merced
05/19–05/20	President, UC Merced Graduate Student Association
05/19–01/20	Graduate student member, Transportation and Parking Services Committee, UC Merced
08/17–05/18	Mentor, Women in STEM (W-STEM) Mentorship Program, UC Merced
04/2017	Volunteer, Center for Cellular and Biomolecular Machines (CCBM) outreach event (Burbank/Peterson elementary schools' visit to UC Merced campus)
08/17–05/20	Treasurer, Merced Indian Graduate Student Association
2017	Co-organiser and website lead, Merced March for Science
08/16–05/17	Member and website lead, BiotaQ: The UC Merced Science Collaborative
08/10–05/15	Editor, *Sopanam* (College magazine, IISER TVM)

Preface

I have read exactly one dissertation which had a preface, and part of me wonders if the decision to tack one on is the result of a need to overshare. But this has been a long endeavour that often seemed impossible, and oversharing or not, I cannot shake the feeling that using this space to provide more context about this work will serve to make it more than the sum of its parts.

I *loved* being a graduate student (the stress and the mental health issues not withstanding), so much so that, if I could simply continue working with Ajay and Justin, I would gladly do so. And I think it has largely to do with how they supported me in exactly the ways I needed and let me chart the path of my PhD. Of course, that resulted in this 100+ page work which, to many, may seem like something without a coherent thread or narrative structure. After all, what do bacteria, human babies and predator-prey systems have in common, right? I think I do a decent of job of providing a technical justification of why this work serves as a cohesive PhD dissertation, but this is where I want to tell you exactly what this means to me and what I see when I look at it.

I have broad interests, as far as my scientific inclinations go: I am just as curious about the dynamics of ecological communities as I am about the way growing up in a multi-lingual household shapes a child's language learning or about the elements that make you want to keep reading a book at 3 am, eyes drooping in exhaustion. And so, over the course of my PhD, I kept collecting projects, and somehow, magically, miraculously, both Ajay and Justin were on-board with my hodge-podge approach to doing a PhD. They did ask me every now and then what the overarching story was, but they also helped me tease it out as I panicked because I did not have an answer other than *this seemed cool and I wanted to do it.*

Now, after nearly six years of dodging that question, I can see it so clearly, how it all fits together. This dissertation is about how organisms forage for resources and how the way

they interact with each other and their environment while going about it shapes the search process. It may seem like a tenuous thread, but this work presents manifestations of one of the most pervasive processes in the world in three wildly different systems, and I cannot help but marvel at how the answer to what was the most terrifying question of my PhD was so deceptively, beautifully–even mundanely–simple. It is all about something that is a forgone conclusion in a world filled with so many things, living and otherwise: organisms interact with each other and (incredible) things happen because of that, and *I* get to study those things in the context of foraging and tease out the connections.

Looking back, I can think of this as nothing but a work of sheer joy, even though I remember the many, many bad mental health days with excruciating clarity. But I think that's the point, that the wonderful aspects shine bright enough to obscure the not-so-wonderful ones. And so, I think this is how I would like this whole doing research thing to go, a random walk in the space of all the curiosity-piquing questions, and I get to work on what I find. Being able to do it this way so far has been a wonderfully liberating experience, and after all, following your curiosity is what being a scientist is really about, right?

But, enough with the rambling. I do feel like I have gotten my point across, albeit circuitously. I'll leave you with this: I feel like I might have done a shoddy job of putting this dissertation together, but I am (somehow) also incredibly proud of what I have been able to do here, and I am enormously grateful to the people (and especially Ajay and Justin) who made it possible. I hope you like this dissertation at least a fraction of how much I like it, and that you are able to see the subtle, unifying thread that brings it all together.

Abstract

While foraging phenomena have historically been a focus of ecology, they can be described as search processes carried out by one or multiple agents in a landscape characterised by a resource distribution. As a result, tools from statistical physics can be used to study foraging phenomena. The integration of statistical physics as a set of tools to both describe patterns observed in foraging processes and understand the underlying mechanisms that give rise to these patterns has enabled the development of robust theory as well as the expansion of the scope of foraging to include applications in intracellular processes, human cognitive behaviour, and even robotics.

Treating foraging processes as undertakings by a single foraging agent in a resource environment has laid the foundation for a significant amount of our understanding of foraging. However, these approaches fail to capture the complexities of real foraging phenomena which arise from the fact that foraging agents do not exist independent of their environment. The logical first order step to rectify this is to consider foraging processes carried out by multiple interacting agents. This dissertation describes complex, emergent phenomena in three different systems resulting from interactions between multiple foraging agents.

First, I present a minimal, one-dimensional analytical model that describes emergent pattern formation in a community of phototactic bacteria moving towards a directional light source. During phototaxis—essentially, foraging for light—interactions between individual bacteria result in emergent community-level spatial organisation, leading to the formation of finger-like projections at the propagating front. We developed a one-dimensional analytical model which describes the dynamics of this pattern formation and predicts critical parameters that limit finger formation. This model also predicts the loss of instabilities in mutant phenotypes lacking a key photoreceptor.

Second, I present a framework to use foraging theory to investigate the dynamics of vocal interactions between human infants and their adult caregivers [2]. By analysing day-long recordings of infant and caregiver vocalisations in naturalistic settings over the infants' first year, we demonstrated support for the hypothesis that infants and their adult caregivers are foraging in an acoustic space for sounds that have social value. Our findings also provide evidence that vocal interactions between infants and caregivers modify these foraging patterns.

Finally, I present a computational model to study the diversity of foraging strategies employed by terrestrial carnivore predators. Mammalian carnivores' foraging strategies include hunting, scavenging, and kleptoparasitism (stealing). However, despite the prevalence of literature on predator-prey systems, the factors that result in the deployment of these strategies and their effects on predator-prey systems are not well understood. In this study, we use an energetics approach to investigate how a focal predator's interactions with potential prey and other predators constrain the use of these strategies. Our results predict the dependence of predator foraging strategy on predator energetics as well as the body sizes of the focal predator, prey, and potential competitors. In particular, our predictions for the boundaries between hunting and alternative foraging strategies (scavenging and stealing) show remarkable agreement with observational data. By employing dimensional reduction, we are also able to accurately describe the phase transition from a state where the focal predator relies predominantly on hunting to a state where the focal predator largely relies on scavenging and stealing.

[2]This chapter is a reproduction of [1]

Contents

List of Tables

List of Figures

Chapter 1

Introduction

We are almost always searching for something, and not just in the more philosophical sense of searching-for-happiness or searching-for-the-meaning-of-life. We search for food, clothes, and a number of other things in stores (or increasingly, on Amazon.com); we search for that one gas station where gas prices are lower than every other gas station in town; and with the internet at our disposal, we are almost constantly searching for information. And before, for a long time in our history, we were searching for food and for shelter from the elements on foot. Suffice it to say, search processes have been—and still are—an integral part of how we go about our lives.

More importantly, search processes that are not carried out by us shape our lives and the world around us in very significant ways. About 75% of global agriculture crop, to different degrees, depend on animal pollination, which is usually a consequence of pollinators foraging and reaching inside flowers for nectar [2,3]. Insectivores such as bats, lizards, and frogs help control pests as a direct consequence of their search for food [4]. Even inside our bodies, vital cellular functions such as DNA transcription and repair depend on DNA-binding proteins' search for specific target sites [5].

Naturally, search processes—especially those that occur in the natural world—have enjoyed enduring scientific interest, and are formally referred to as *foraging* processes. Historically, foraging has been defined as the collection of processes that organisms engage in to obtain the resources necessary for survival and/or reproduction [6]. Naturally, foraging studies predominantly fall (or more accurately, *fell*) within the purview of ecology. However, in its most bare-bones form, foraging is simply a search process carried out by one or multiple agents in some landscape characterised by a resource distribution, and hence, can

be modelled as a random search [7].

Biologists have long recognised the utility of statistical physics as a tool to analyse the patterns in the search processes of organisms. However, the integration of physics as a theoretical framework capable of contributing to the understanding of foraging processes is more recent [8]. Arguably, the seminal work in this context is reported in a 1996 paper by Viswanathan et al. [9] where the authors found that the probability distribution of the flight times of the wandering albatross (*Diomedea exulans*) followed a heavy-tailed power law distribution. This kind of search pattern is known as a Lévy flight, and has been shown to be an optimal foraging strategy when resources are sparsely distributed [7, 10, 11].

Over time, the scope of foraging studies has expanded beyond (macro)ecology, to include foraging processes of microorganisms [12] and beyond. It is now known that foraging phenomena are found in human cognitive behaviour [13–15], online social networks [16], information spaces [7, 17], etc. Our understanding of physics has played an important role in this expansion, by identifying patterns arising in foraging processes and building a generalised theoretical framework to understand these patterns [7, 8, 15].

As mentioned previously, foraging processes may be carried out by one or more agents in search of a resource. For example, solitary predators such as tigers (*Panthera tigris*) [18] and leopards (*Panthera pardus*) [19] hunt by themselves. Other examples of solitary foragers include the eastern carpenter bee (*Xyloeopa virginiea*) [20], the northern pike (*Esox lucius*) [21], white-tailed trogon (*Trogon viridis*) [22], etc. Group foraging, on the other hand, involves multiple agents collectively—and often collaboratively—searching for resources [23].

The relatively simple case of individual foragers searching for resources has been used to develop optimal foraging theory (OFT), which posits that some currency is *optimised* during foraging. This currency can be the rate of energy intake by the forager, the total energy intake by the forager, or the forager's lifetime reproductive success [24–26]. Since its introduction in 1966, OFT has received strong support as well as vehement criticism, and in general, has only been successful in adequately describing simple realisations of the foraging process [27, 28].

The shortcomings of OFT can (perhaps) be best summarised thus: organisms do not exist independent of each other, and a realistic understanding of complex phenomena like foraging requires taking several complicating factors into account [28–30]. One such complicating factor is the fact that often, foraging processes are undertaken by more than one foraging agent.

Sometimes, this takes the form of organisms foraging collaboratively in groups, since the interactions facilitated by virtue of working together are often helpful in increasing the average rate of resource consumption [31, 32]. Often, the splitting of foraging costs and benefits between agents in a group foraging context confers a collective advantage to the group that individual foragers cannot achieve. To put it simply, when organisms forage in groups, "the food discoveries of a few can lead to the feeding of many" [33]. In other cases, multi-agent foraging simply amounts to more than one agent foraging—and competing—for the same (often finite) resource. For example, lions (*Panthera leo*) and spotted hyenas (*Crocuta crocuta*) show significant overlap in their prey preferences [34]. When multiple agents engage in foraging, collaboratively or otherwise, they may interact directly [35, 36] or indirectly (eg. by modifying the target resource distribution) [37] with each other. These interactions give rise to complex behaviours which lead to optimal group sizes, the emergence of scroungers who steal resources from other foragers, temporal and spatial patterns in the distributions and behaviours of the foragers and the target resource, etc. [31, 33, 38–40].

Equally important to consider is the resource environment the agents are foraging in: the distribution, abundance, and accessibility of the target resource play an important role in shaping foraging patterns. For example, Lévy flights have been shown to be optimal foraging strategies for sparse resource distributions that have no limitations on the number of times a site can be visited by a forager [10]. On the other hand, foragers in resource-dense regions tend to use smaller steps and paths with large tortuosities (or twistedness) [41].

In this text, I present my work on interactive multi-agent foraging phenomena in three systems: phototactic bacterial communities, human infants learning language, and predator-prey systems with multiple predator foraging strategies. While these may seem disparate, the work presented here is linked by the perspective of how interactions between multiple foraging agents and their resource environment modify foraging phenomena across different scales. More importantly, this dissertation is a demonstration of both the diversity of foraging phenomena and the versatility of using tools from physics in understanding them.

In Chapter 2, we investigate collective foraging and dynamic pattern formation in communities of phototactic cyanobacteria. *Synechocystis* sp. PCC6803 is a photosynthetic, photactic freshwater cyanobacterium that moves directionally in response to a light source by gliding [42]. During phototaxis—which, in the case of phototactic bacteria, amounts to foraging for a nutritional source—interactions between individual bacteria result in the aggregation of bacteria in a thin, dense strip at the propagating front followed by emergent

spatial organisation in the form of finger-like projections [39]. We propose a one-dimensional analytical model which describes the dynamics of finger formation, predicts critical parameters that limit finger formation and predicts the observed loss of instabilities in *taxD1* mutants, which lack an important photoreceptor. Predictions from this minimal, analytically solvable one-dimensional model can be tested against results from existing 2-dimensional computational models [39], as well as experimental data obtained from phototaxis assays.

In Chapter 3, we leverage our understanding of foraging theory to frame vocal interactions between human infants and their adult caregivers as a foraging process [1]. Vocal learning observed in human infancy results from a combination of social rewards, imitation of adult speech, and exploratory processes [43–49]. We investigated the hypothesis that infants and adult caregivers are foraging in an acoustic space for sounds that have social value, using day-long recordings of vocalisations in a naturalistic setting over the infants' first year. Our findings are consistent with the hypothesis that infants forage vocally for social input and that vocal interactions between infants and caregivers modify these foraging patterns. This study is a novel application of foraging theory to characterise infant-caregiver vocal interactions by assessing vocal exploration in terms of patterns of movement in an acoustic space, which will allow this domain of behaviour to be compared to other foraging behaviours.

In Chapter 4, we build a computational model to understand how a focal predator's interactions with prey and other predators result in different predator foraging strategies. The foraging behaviours that organisms employ can be diverse and have tangible energetic costs. In the case of mammalian carnivores (predators) searching for food, these behaviours include hunting, scavenging, and stealing (kleptoparasitism). While there is a wealth of literature—both theoretical and observational—on predator-prey systems, they largely focus on hunting and its consequences, and our understanding of scavenging and kleptiparasitism is sorely lacking. We investigate how pairwise interactions between a focal predator, prey, and potential *competitors* result in the deployment of different predator foraging strategies. Dimensional reduction of our results reveal a behavioural switch—a phase transition, if you will—between hunting and alternative modes of predation (scavenging and stealing) as a function of prey mass. In addition, these results show remarkable agreement with observational data, and in particular, predict the boundaries between foraging strategies exceptionally well.

Finally, in Chapter 5, I summarise this work and present conclusions.

Chapter 2

A One-Dimensional Model for the Initiation of Fingering Instabilities During Cyanobacterial Phototaxis

2.1 Introduction

Phototaxis is movement in response to light. A number of organisms across taxa are known to be phototacic, including zooplankton, cyanobacteria, jellyfish, and insects [42, 50–52]. When the organism in question moves towards the light source, the movement is called positive phototaxis, while movement away from the light source is called negative phototaxis.

This work focuses on the cyanobacterium *Synechocystis* strain PCC 6803 (referred to as *Synechocystis* henceforth), which is a well-studied photosynthetic and phototactic cyanobacterium [42]. *Synechocystis* is a unicellular cyanobacterium found in freshwater sources [53]. It can undergo phototrophic growth through photosynthesis in the presence of light as well as heterotrophic growth through glycolysis and oxidative phosphorylation in the absence of light [54, 55].

These cyanobacteria secrete an extracellular polymeric substance (EPS) which enhances the mobility of the bacteria. Thus, by secreting EPS, individual bacteria are able to not only enhance their own motility but also that of other bacteria in the vicinity, resulting in interactions between bacteria [39]. These interactions between individual bacteria result in emergent pattern formation in communities of *Synechocystis* during positive phototaxis [39].

This pattern formation can be best described as finger-like projections that emerge from an initially flat propagating bacterial front (Fig. 2.1). In a typical phototaxis assay, cells spotted onto a moist, low concentration agarose surface are placed in a directional light source, and the finger-like formations emerge after about a day of the system evolving [39, 56].

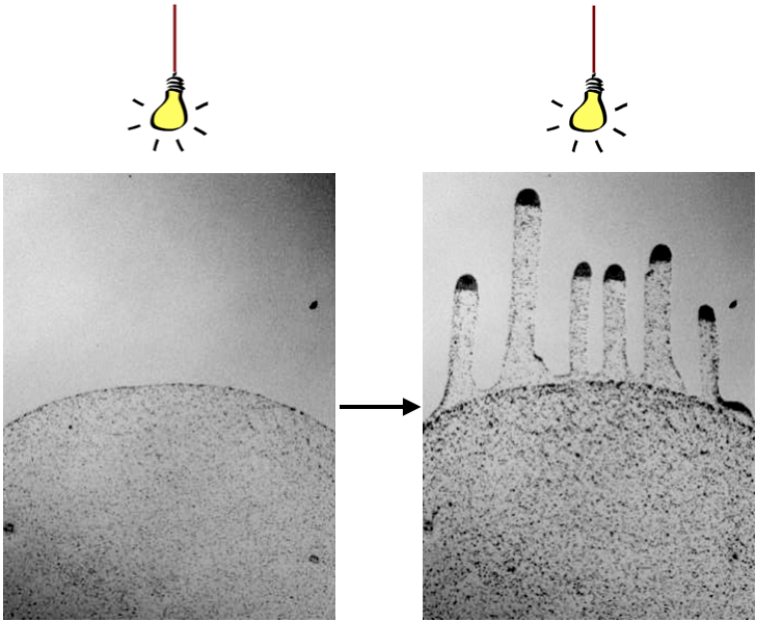

Figure 2.1: **Finger-like projections emerge from an initially homogenous drop of *Synechocystis* cells during the course of phototactic motion.** The light source is placed above the bacterial drop as shown, and the fingers move towards the light source. Images from [39, 56]

This pattern formation is the result of complex signal transduction pathways which involves detecting the directional light source and local EPS concentrations, and transducing them into appropriate motility responses [57, 58]. Several components underpinning this motility response have been identified using genetic and molecular analysis [59–62]. In particular, it is now known that Type IV pili (TFP) and TaxD1, a cyanobacteriochrome photoreceptor, play important roles in mediating the observed community behaviour in *Synechocystis* in response to a directional light source [56, 59, 63, 64]. TFP is necessary for motion on a surface and adhesion, and mutants that lack TFP are completely non-motile [63]. TaxD1 promotes positive phototaxis and inhibits sideways motion (motion

perpendicular to the direction of the light source) [59, 64]. In fact, mutants with TaxD1 knocked off show increased sideways motion as well as negative phototaxis, in addition to exhibiting a loss of finger formation, indicating that while TaxD1 is necessary for phototaxis, there may be other photoreceptors involved in mediating the phototactic response [56,59,64].

In light of the complexity of the underlying mechanism of emergent pattern formation in this system, biophysical models can shed light on mechanistic aspects of of this phenomena. Previous work has resulted in a two-dimensional computational model based on experimental observations, and successfully reproduces the finger formation [39]. However, this model is not analytically solvable, and only numerical solutions exist which limits its predictive ability.

In this work, we build a minimal, analytically solvable one-dimensional model to describe this phenomenon. This is possible because over the course of the phototactic response, the bacteria aggregate in a high-density band at the leading edge, enabling a one-dimensional approximation of the system.

2.2 Theory

2.2.1 A one-dimensional theory for the initiation of fingering instabilities during phototaxis

Since we are interested in the fingering transition of the inoculation boundary due to phototaxis, we model the height of the migrating front at time t relative to a flat surface as a one-dimensional curve $h(s,t)$, where s represents length along the contour, with cellular concentration $c(s,t)$ along the front. The curve begins as a horiztonal line at $t = 0$. When fluctuations transform the shape of the front away from being flat, the flux of cells along the front has contributions from the phototactic bias and from random diffusion:

$$J = D_s \left(f sin\theta c - \frac{\partial c}{\partial s} \right) \tag{2.1}$$

where f is an effective force representing the phototactic bias, D_s is the diffusion constant along h, and θ is the angle between the horizontal and the tangent to the curve at s. We assume the phototactic force to always be directed perpendicular to the initial shape of the front. In addition, we assume that θ is small, such that $sin\theta = \partial h/\partial s$.

Assuming no gain or loss in cell density,

$$\frac{\partial c}{\partial t} = -\nabla J = \frac{\partial}{\partial s}\left(D_s \frac{\partial c}{\partial s} - D_s f \frac{\partial h}{\partial s} c\right) \tag{2.2}$$

We consider short times such that h displays only small deviations from flatness and assume that all displacements are normal to the interface, giving

$$\frac{\partial h}{\partial t} = f\, D_p \,+\, \frac{\sigma}{\eta} d\, \nabla^2 h \tag{2.3}$$

where D_p is the diffusion constant perpendicular to the interface and is not necessarily equivalent to D_s (for instance due to heterogeneous distribution of TFP across the cell surface) and d is the thickness of the EPS layer. In Eq. 2.3, we have ignored a factor of $cos\theta$ in the first term based on our assumption of small θ. In the second term, $\sigma\nabla^2 h$ is the force resulting from the surface tension of the EPS interface, producing an interface velocity according to the effective drag η.

To model the density-dependent production of EPS, we assume D_s and D_p are proportional to the local cell density c, equivalent to assuming that the EPS concentration equilibrates very quickly compared to cellular movement:

$$D_s = \alpha\, c = \beta\, D_p \tag{2.4}$$

From Eqs. 2.3 and 2.4, we arrive at

$$\frac{\partial h}{\partial t} = \bar{f}\, c + \,\bar{\sigma}\, \nabla^2 h \tag{2.5}$$

where $\bar{f} = f\,\alpha/\beta$ and $\bar{\sigma} = \sigma d/\eta$. Solving for c,

$$c = \bar{f}^{-1}\left(\frac{\partial h}{\partial t} - \bar{\sigma}\, \nabla^2 h\right) \tag{2.6}$$

Substituting Eq. 2.6 into Eq. 2.2,

$$\frac{\partial^2 h}{\partial t^2} - \bar{\sigma}\, \frac{\partial^3 h}{\partial t \partial s^2} = -\frac{\partial}{\partial s}\left(\beta\, \frac{\partial h}{\partial s}\, \Gamma^2 - \frac{\alpha}{\bar{f}}\, \Gamma\, \frac{\partial \Gamma}{\partial s}\right) \tag{2.7}$$

where $\Gamma = \partial h/\partial t - \bar{\sigma}\, \partial^2 h/\partial s^2$. For uniform initial conditions, $\partial h/\partial s = 0$, yielding $h = v_0 t$, where $v_0 = \bar{f} c_0$ with c_0 as the initial uniform concentration.

2.2.2 Linear stability analysis of fingering instabilities

To determine the onset of instabilities, one can perform a linear stability analysis around the uniform solution as

$$h_p(s,t) = v_0 t + \epsilon \hat{h}(s,t) \tag{2.8}$$

Substituting this expression into Eq. 2.7 and retaining terms linear in ϵ yields

$$\frac{\partial^2 \hat{h}}{\partial t^2} - \bar{\sigma}\,\frac{\partial^3 \hat{h}}{\partial t \partial s^2} = \frac{-\partial^2}{\partial s^2}\left(\beta\, v_0^2\, \hat{h} - \frac{\alpha}{\bar{f}}\, v_0 \hat{\Gamma}\right) \tag{2.9}$$

where $\hat{\Gamma} = \partial \hat{h}/\partial t - \bar{\sigma}\,\partial^2 \hat{h}/\partial s^2$. Substituting solutions of the form $\hat{h} \sim e^{\lambda t} e^{iks}$ into Eq. 2.9, we obtain a quadratic equation for λ:

$$\lambda^2 + \lambda\left(\bar{\sigma} + D_0\right) k^2 + D_0\,\bar{\sigma}\,k^4 - \frac{f^2\, D_0^2 k^2}{\beta} = 0 \tag{2.10}$$

using $D_0 = \alpha\, c_0$ and $\beta\, v_0^2 = f^2\, D_0^2/\beta$. Solving for λ gives

$$\lambda = \frac{k^2\big(D_0 + \bar{\sigma}\big)}{2}\left[-1 \pm \sqrt{1 + \gamma\left(\frac{k_c^2}{k^2} - 1\right)}\right] \tag{2.11}$$

where $k_c^2 = f^2 D_0/(\beta\bar{\sigma})$ and $\gamma = 4 D_0 \bar{\sigma}\big(D_0 + \bar{\sigma}\big)^{-2}$. Note that the determinant is always positive (the maximum value of γ is 1), and hence λ is always real. For a detailed derivation of Eq. 2.11, see Section A.1.

The dependence of λ on k is non-monotonic (see Fig. 2.3a): the growth rates of unstable modes increase from $\lambda \to 0$ when $k \to 0$, till it attains a maximum value, $\lambda_{\max}$, corresponding to the fastest growing unstable wavelength, characterised by its wave number $k_{\max}$. The wavelength corresponding to $k_{\max}$ is the wavelength of the dominant instability in the system, provided that instabilities have not merged together (Fig. 2.2a). To determine $k_{\max}$, we set $\partial\lambda/\partial k = 0$ and obtain, after some algebra (see Section A.2),

$$k_{\max} = k_c\big(2\big(1 + \gamma^{-0.5}\big)\big)^{-0.5} \tag{2.12}$$

Some more algebra (see A.2) provides a simplified expression for $k_{\max}$ in terms of fewer derived quantities:

$$k_{\max} = \frac{f D_0^{3/4}}{\bar{\sigma}^{1/4}\sqrt{\beta}\big(\sqrt{D_0} + \sqrt{\bar{\sigma}}\big)} \tag{2.13}$$

By setting $k = k_{\max}$ in 2.11, we can also derive an expression for the growth rate of the fastest growing mode, $\lambda_{\max}$ (for a detailed derivation, see Section A.2):

$$\lambda_{\max} = \frac{f^2 D_0^2}{\beta\left(\sqrt{D_0} + \sqrt{\bar{\sigma}}\right)^2} \tag{2.14}$$

The growth rate of modes greater than $k_{\max}$ decreases till it falls to 0, corresponding to the critical wave number, k_c. That is, as $k \to k_c$, $\lambda \to 0$. Alternatively, all modes with $k > k_c$ decay instead of growing, and hence, do not contribute to finger formation. This critical wave number value, as described earlier in this chapter, is given by

$$k_c = f\sqrt{\frac{D_0}{\bar{\sigma}\beta}} \tag{2.15}$$

Finally, we note that the formation of observable fingers is limited by the physical dimension of the system, L. To observe an unstable mode requires that the system size, L, be larger than the shortest unstable mode, i.e. $k_c > 2\pi/L$. This imposes a lower limit on the phototactic bias for (observable) finger formation (see Fig. 2.2b) such that

$$f > \frac{2\pi}{L}\sqrt{\frac{\beta\bar{\sigma}}{D_0}} \equiv f_c$$

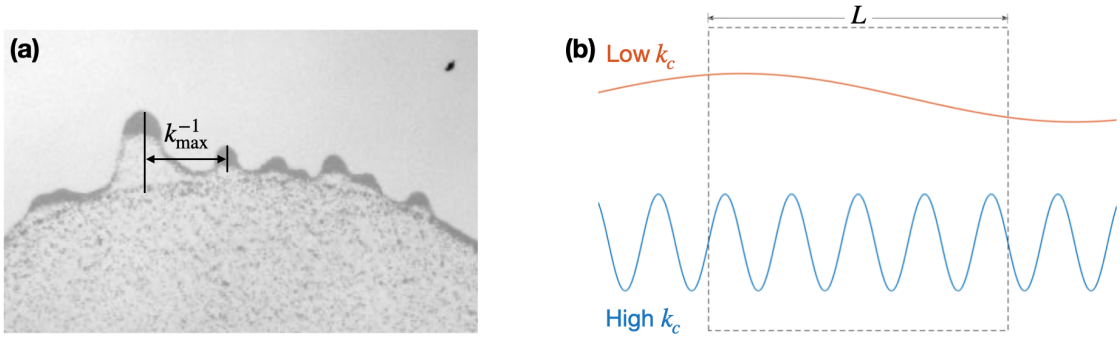

Figure 2.2: **Explanatory schematics for wavelengths associated with the fastest growing mode, $k_{\max}$ and critical mode, k_c.** (a) The fastest growing mode, $k_{\max}$ is associated with the wavelength (or the distance between) the dominant fingers in the system, provided multiple fingers have not merged together. (b) shows how the system size limits pattern formation in the. Low k_c values (red curve) correspond to large wavelengths. If this wavelength is larger than the system size L, finger formation cannot be observed. High k_c values (blue curve) correspond to small wavelengths. As a result, finger formation will be observable for a larger range of system sizes (image in (a) from [56]).

2.2.3 Alternative cases

In the analysis above, with D_s, $D_p \propto c$, we found that the condition for instability was $k^2 < k_c^2$ (Eq. 2.15). Several other scenarios can be readily studied with small modifications; in each case Eq. 2.12 for $k_{\max}$ remains the same, while the definition of γ changes. If we ignore the cell density dependence of D_s, focusing only on cell density/EPS effects for motion out of the inoculation (which is primarily dictated by D_p), γ becomes $4D_s\bar{\sigma}/(\bar{\sigma}+D_s)^2$. Instabilities are still possible, but emerge below a critical wave vector $k^2 < k_c^2 D_s/D_0$ as long as the system size is above $\frac{2\pi}{f}\sqrt{\frac{\bar{\sigma}\beta}{D_s}}$.

In the trivial case in which both D_s and D_p are constants independent of cell density, no instabilities can occur: Eq. 2.11 becomes $\lambda = -k^2\sigma/\eta$, which is always negative and hence no unstable modes exist.

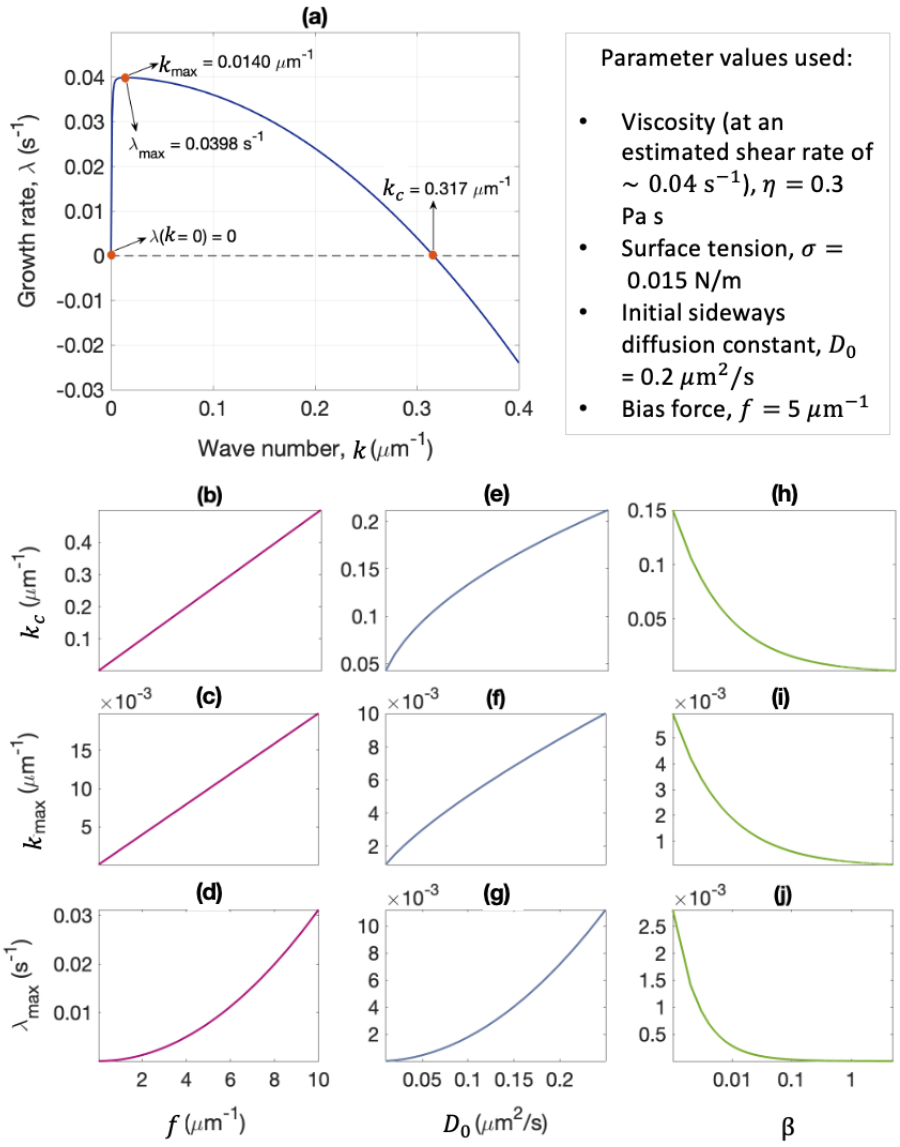

Figure 2.3: **A summary of (testable) predictions from the one-dimensional model.** (a) The growth rate of instabilities, λ, as a function of the wave number, k, is shown. This functional form was computed by estimating the surface tension of EPS, $\sigma \sim 0.010.06$ N/m [65–67], and the viscosity of EPS, $\eta \sim$ 10-300 mN s/m2 [68, 69]. D_0, d, and f were estimated from [39]. Exact values used are shown in the textbox on the top right. k_c, $k_{\max}$, and $\lambda_{\max}$ for these parameter values are indicated. (b-d) The dependence of the critical mode, k_c, on the bias force, f, the initial perpendicular diffusion constant, D_0, as well as the ratio of the strength of sideways diffusion and perpendicular diffusion, β are shown. (e-g) The dependence of the fastest growing mode, $k_{\max}$, on f, D_0, and β are shown. (h-j) The dependence of the growth rate of the fastest growing mode, $\lambda_{\max}$, on f, D_0, and β are shown. For (b-j), all parameter values are the same as for (a) except the parameter being varied.

2.2.4 Testable predictions

Equations 2.13–2.15 summarise three key variables which can be used to characterise the system. $k_{\max}$ is the wave number associated with the fastest growing instability mode, which provides an estimate for the wavelength of the dominant finger (Fig. 2.2a), while $\lambda_{\max}$ provides an estimate for how fast this instability grows. Together, these two quantities set an upper limit for the size of the finger-like projections in the system—instability wavelengths larger than the wavelength associated with $k_{\max}$ cannot be easily observed in the system (even though unstable wavelengths larger than $k_{\max}$ are present in the system, due to the high growth rate of the $k_{\max}$ mode, it will dominate all other wavelengths)—as well as a timescale for finger formation. k_c, the critical wave number, is associated with the smallest instability wavelength present in the system, and hence, sets a lower bound for observable finger wavelength.

Equations 2.13–2.15 also describe how these three system descriptors depend on several tunable parameters, which allows for testing the validity of the one-dimensional model outlined here against results from experiments conducted in the laboratory, *in silico* [39], or both. For example, the wave number of the fastest growing mode and the critical wave number are expected to increase linearly with the phototactic bias force (Fig. 2.3b, c), i.e., with increase in the phototactic bias force, smaller finger widths contribute to the observed pattern formation *and* the wavelength of the fastest growing mode becomes smaller. These predictions have already been shown to be qualitatively true through *in silico* experiments [39]. The one-dimensional model presented here provides analytical expressions for these (and other) dependences as summarised in equations 2.13–2.15 and Fig. 2.3b-j, which can be tested using information from a time series of the advancing bacterial front.

In particular, our model describes the dependence of $k_{\max}$, $\lambda_{\max}$ an k_c on β, a ratio which describes the strength of bacterial diffusion along the one-dimensional curve (or alternatively, in the full two-dimensional context, diffusion in the X direction) with respect to that of diffusion perpendicular to the one-dimensional curve (Y direction and directed towards the light source). While this is not relevant for wild type bacteria, the differential diffusion strengths play a crucial role in the loss of finger formation in *taxd1* mutants, which we will address next.

2.2.5 Modification of the model to explain *taxD1* cells

taxD1 cells exhibit increased diffusion in the direction perpendicular to the light source, as well as negative phototaxis (movement away from the light). While negative phototaxis can simply be modeled as a reversal of the direction of the bias force, we aimed to use our model to interpret the changes in the front morphology in this mutant. We considered an increase in D_s, or equivalently an increase in β, which results in a decrease in k_c and an increase in f_c. Thus, instability sets in only at longer wavelengths, potentially explaining the absence of fingering instabilities in *taxD1* inoculations, though sufficiently large systems should have the potential to generate instabilities.

2.3 Conclusion

Here, we have constructed an analytically solvable one-dimensional model based on minimal assumptions to describe the formation of finger-like patterns in wild type *Synechocystis* communities, and the loss of finger formation in *taxD1* mutants. By virtue of being analytically solvable, this model provides predictions (Eqs. 2.11, 2.13-2.15) about the effect of tunable system parameters on pattern formation, which can be tested *in vivo* and/or *in silico*, while also enabling the exploration of larger swaths of the system parameter space than would be possible experimentally. However, not all predictions presented here may be as readily testable *in vivo* as they are *in silico*. For instance, tuning the phototactic bias, f, will likely prove to be much easier in laboratory experiments than altering D_0, since it is a rather more emergent parameter than f. Similarly, while the effect of β can be easily tested by the addition of an anisotropic diffusion tensor in the two-dimensional simulations [56], the measurement of said quantity may prove harder in experiments with *taxd1* mutants. Nevertheless, the one-dimensional model provides a powerful analytical handle on this system and opens the doors to higher order explorations.

Chapter 3

Exploratory Dynamics of Vocal Foraging During Infant-Caregiver Communication

3.1 Introduction

3.1.1 Infant vocal development

Human infants show massive growth in vocalising abilities during their first year [70–74]. Cries and short, quiet sounds dominate early vocalisations. By about three months, infants demonstrate a much wider range of vocalisation types, varying pitch, amplitude, and other phonatory characteristics, as well as beginning to produce primitive consonant-vowel articulations. By 7 months, infants begin producing well-timed adult-like consonant-vowel alternations. This expansion in repertoire lays a foundation for later speech and other vocal communication production [71, 75]. For instance, there is continuity between the sounds produced during prelinguistic stages of vocal development and those sounds that make up infants' first words [76].

Although anatomical changes and neuromaturation may account for some of the changes during the first year, even newborn infant vocal tracts are capable of producing a very wide range of sounds, and the dramatic changes in infants' vocalisations over the first year are believed to be primarily due to learning [77–79]. Computational models of infant vocal learning have demonstrated that some combination of exploratory processes, social

or intrinsic rewards, and imitation of adult speech input can result in the vocal learning we see in human infancy [43–49]. Infant vocal learning may be viewed as a process that combines variation and selection, resulting in the evolution of a more adult-like repertoire of sound types. From this perspective, it is important to know more about the dynamics of the generation of varied sounds and the factors that contribute to the selection of some of those sounds over others.

A large body of research with human infants provides strong support for the role of social responses in shaping infant vocal learning at multiple timescales. Adult responses are said to be contingent on infant vocalisation type if a response is more likely for some types of vocalisations than for other types. Adult responses contingent on speech-related infant vocalisations result in more frequent speech-related and adult-like vocalisations by the infant during the seconds and minutes following the adult response [79–84]. At the longer timescales of months and years, differences in adult responses predict later communication abilities [85–87]. Differences in feedback loops wherein infant vocalisations generate adult responses that in turn impact future infant vocalisations appear to underlie some of the differences in speech development that are observed on average across socioeconomic and clinical groups [83, 88, 89].

Less is known about exploratory dynamics in the vocal domain over the course of the day. Research has documented temporal clumping in the occurrence of different prelinguistic vocalisation types, termed "session effects" because the effect is that one recording session can provide a very different view of an infant's vocal repertoire than one made a few minutes later [90]. Fractal-like clustering of infant vocalisation events in time has been found within day-long audio recordings, with bouts of infant vocalisations clustered within larger episodes of high-volubility with those in turn clustered at even longer timescales [91]. The same study also found that the adult vocalisations to which infants were exposed were also hierarchically clustered in time, with the degree of hierarchical clustering being coordinated between infants and adults.

The finding of coordinated hierarchical clustering in the timing of adult vocalisations, and the fact that adult responses appear to play such an important role in infant communication development, highlight adults as active agents in the infant vocal exploration process. Indeed, a large volume of prior research has documented how adults alter the timing, acoustic properties, and content of their vocalisations when they are directed to an infant [92–98]. While less research has focused on how future adult behaviour is influenced

by infant responses to adult vocalisations, we do know that infant vocalisations tend to follow adult vocalisations during naturalistic interactions [99–101] and that infant responses are sensitive to differences in adult behaviours [102]. These findings taken together suggest that it is appropriate to think not only of infants as vocal explorers in search of sounds that have social value but also that adults may also be doing a similar kind of foraging.

3.1.2 The interdisciplinary study of foraging

We propose that infant and adult vocalisations can be viewed in part as a process of exploring various sound types in search of those that tend to yield social responses. This perspective makes it possible to compare infant and adult vocal behaviour to other types of foraging, such as animals foraging in physical space to collect food, identify mates or find safe locations to escape predators. Other non-spatial behaviours, such as adult humans' search in semantic networks during memory retrieval tasks, have already benefited from the application of methods previously used to study animal foraging in space [103–107].

The search patterns of foraging animals have been shown to hold a wealth of information about optimal foraging strategies under various conditions when analyzed in terms of the distribution of the distances (a.k.a. step sizes) traveled in a particular direction between reorientations [108]. For example, Lévy walks are a random search strategy consisting of mostly short steps interspersed with a "heavy tail" of long steps, characterised by steps of length l occurring with a power law probability distribution $P(l) \sim l^{-\alpha}$ with $1 < \alpha \leq 3$. Lévy-like path length distributions have been reported for a number of spatially foraging organisms. Foraging strategies can provide insight into the clustering of resources. For instance, for sparse and patchy resource distributions, heavy tailed random walks such as Lévy walks, have been shown to be preferred [108]. The scaling exponent, α, of power law distributed path lengths is also informative, differing as a function of the patchiness or clustering of resources [109]. In other cases, foraging appears to follow an exponential pattern, $P(l) \sim e^{-\lambda l}$, where the likelihood of a given step size drops off more rapidly than a power law as step size increases, giving rise to behaviour that is better modelled as Brownian movement. The parameter λ can be fit to the data and reflects the rate of exponential drop-off in likelihood as step size increases. Exponential step size distributions are often observed for animals' spatial foraging when resources are plentiful [110]. Other distributions, including lognormal distributions, defined as $P(l) = (l\sigma\sqrt{2\pi})^{-1}e^{-(\ln l-\mu)^2/2\sigma^2}$, represent step size distributions with tail sizes that are intermediate between exponential

and power law. Thus, depending on the distribution of resources in the environment, different kinds of search patterns can be optimal.

The temporal scaling of behaviour, i.e. the relative frequency of short events (or inter-event intervals) compared with longer events or the degree to which events are clustered hierarchically in time, sometimes changes over the course of learning. For example, it has been observed that the temporal scaling of infant eye movements during an audio-visual speech perception task changes with age, correlates with language ability, and differs for infants with autism spectrum disorder [111]. Adults also show changes in the temporal scaling of their eye movements as they learn in a mathematical insight task [112]. It is possible that the temporal scaling of infant and/or adult vocal behaviour change as infants gain increasing mastery over their vocal motor systems and as they learn about the social consequences of their vocal behaviour, and as adults adapt their behaviour to facilitate infant communication development [91]. As the patterns of the grouping of behaviours in time has been found to be closely related to distances moved in space (both physical and cognitive, e.g. semantic [104, 105]), we would also expect changes in temporal spacing of infant and adult vocal episodes to correspond to exploration of acoustic properties.

3.1.3 The present study

Here we study the dynamics of socially-guided vocal exploration over the course of a day by infants and their caregivers from a foraging perspective. We assess exploration in terms of patterns of movement in acoustic space, an approach which has not previously been employed in the study of infant vocalisations and which will allow this domain of behaviour to be compared to other foraging behaviours, such as animal movement in physical space and human search in memory. We ask whether some patterns of exploration are associated with greater degrees of social responding than others, whether infant exploration is modulated by adult responses and vice versa, and whether these patterns change over the course of the first year.

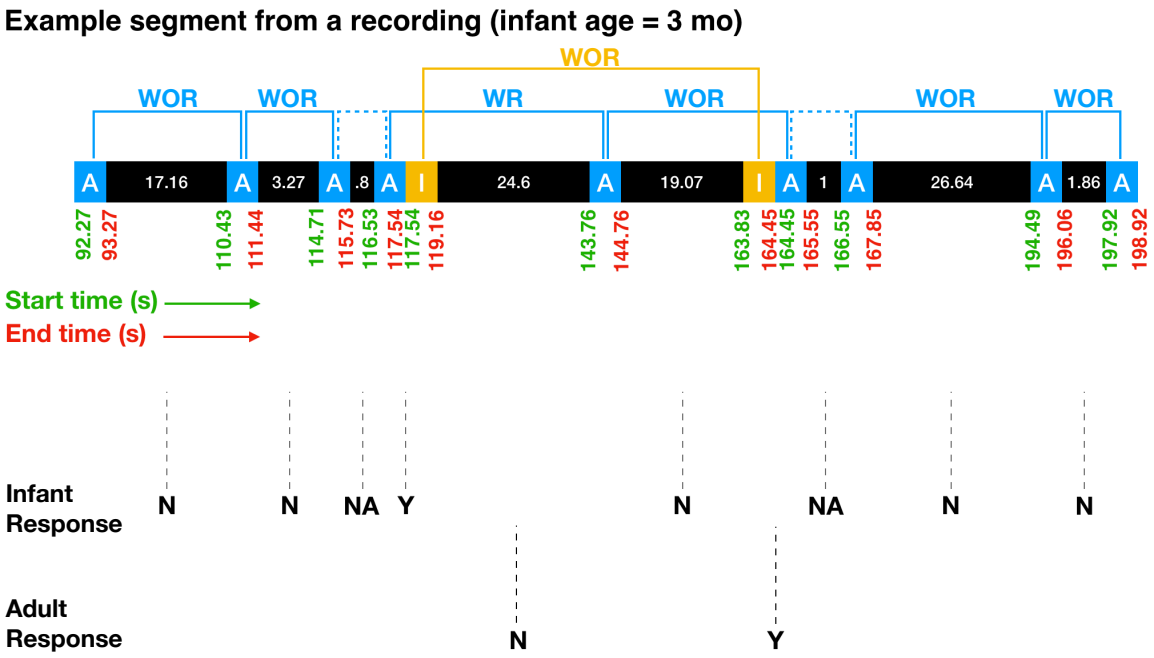

Figure 3.1: **Illustration of the automatic labelling scheme.** Here, we use a small portion from a recording of an infant at age 3 months to illustrate the labelling scheme. The information in the top portion of the figure is provided by the LENA system; inter-vocalization intervals, acoustic features, and response codes are computed automatically based on those LENA-generated speaker labels and the original audio file. Adult vocalisations are in blue boxes while infant vocalisations are in yellow boxes. The start (left, green) and end times (right, red) of each vocalisation are indicated below the corresponding box. Inter-vocalisation intervals (in s) are given in black boxes between subsequent vocalisations, and the length of the box is indicative of the duration of the inter-vocalisation interval. For each infant/adult vocalisation, the receipt of a response (Y/N/NA) is indicated (dotted vertical line). An adult vocalisation is labelled as having been followed by an infant "response" (WR for "with response"; Y in the Infant Response row) if the onset of an infant vocalisation occurs within 1 s following the offset of the adult vocalisation, with no intervening adult vocalisation. The same logic applied for labelling whether adult "responses" followed infant vocalisations. If two adult vocalisations are separated by 1 s or less without an intervening infant vocalisation, then response to the first is NA. Adult vocalisation steps analysed in this study are indicated by the blue lines, with steps following a response labelled WR in blue, and steps following no response labelled WOR (for "without response") in blue. Infant vocalisation steps are indicated by the yellow lines, with steps following a response labelled WR in yellow, and steps following no response labelled WOR in yellow. Dotted lines connecting vocalisations indicate that the corresponding inter-vocalisation interval was not analysed because it was less than or equal to 1 s with no intervening response. Analysed inter-vocalisation intervals were divided into WR and WOR depending on whether the first vocalization in the sequence was followed within 1 s by a "response".

We utilise automatically labelled day-long audio recordings, which have two main advantages. First, they provide very ecologically valid samples of infant and adult behaviour in the range of contexts typically experienced by the infant. Second, they provide long time series with large numbers of events, which is critical in order to characterise the tails of distributions of vocal events. We define acoustic space as the combination of two parameters, the mean pitch and amplitude at the utterance level. While we also recognise that this space does not reflect the complete set of acoustic features relevant to human vocal communication, pitch and amplitude have the advantages of (1) being automatically measurable and (2) being relevant dimensions for most vocalizations produced by infants in the age range studied here as well as for adult vocalizations [90, 93, 94, 101].

3.2 Results

3.2.1 Steps in time and acoustic space

Fig. 3.2 presents an example of visualising infant vocalisation steps in 2D acoustic space and visualising position in each acoustic dimension as a function of the time in the recording at which the vocalisation occurred. Euclidean distance in the acoustic space is defined as $s = \sqrt{\Delta f^2 + \Delta d^2}$, where Δf is the change in mean pitch (log f_0) from the i^{th} to $i + 1^{th}$ vocalisation, and Δd is the corresponding change in mean amplitude. Note that acoustic dimensions are unitless since they have been standardised. Steps for each speaker type were divided into two groups (Fig. 3.1): (1) steps following receipt of a response by the other speaker type (referred to as WR, for 'with response') and (2) steps following a vocalisation that did not receive a response (referred to as WOR, for 'without response').

We found weak positive correlations between distance in 2D acoustic space (i.e., the step size from the i^{th} to $i + 1^{th}$ vocalisation in acoustic space) and the corresponding inter-vocalisation time intervals for both infants (over the 143 recording days, mean $r = 0.10$, mean $p = 0.11$, median $r = 0.09$, median $p = 0.01$, mean $n = 1043.86$, where r is the Pearson correlation coefficient and p is the probability of the null hypothesis) and adults (mean $r = 0.08$, mean $p = 0.06$, median $r = 0.08$, median $p < 0.001$, mean $n = 3223.40$). Note that the number of adult vocalisation events is much larger than that of infant vocalisation events.

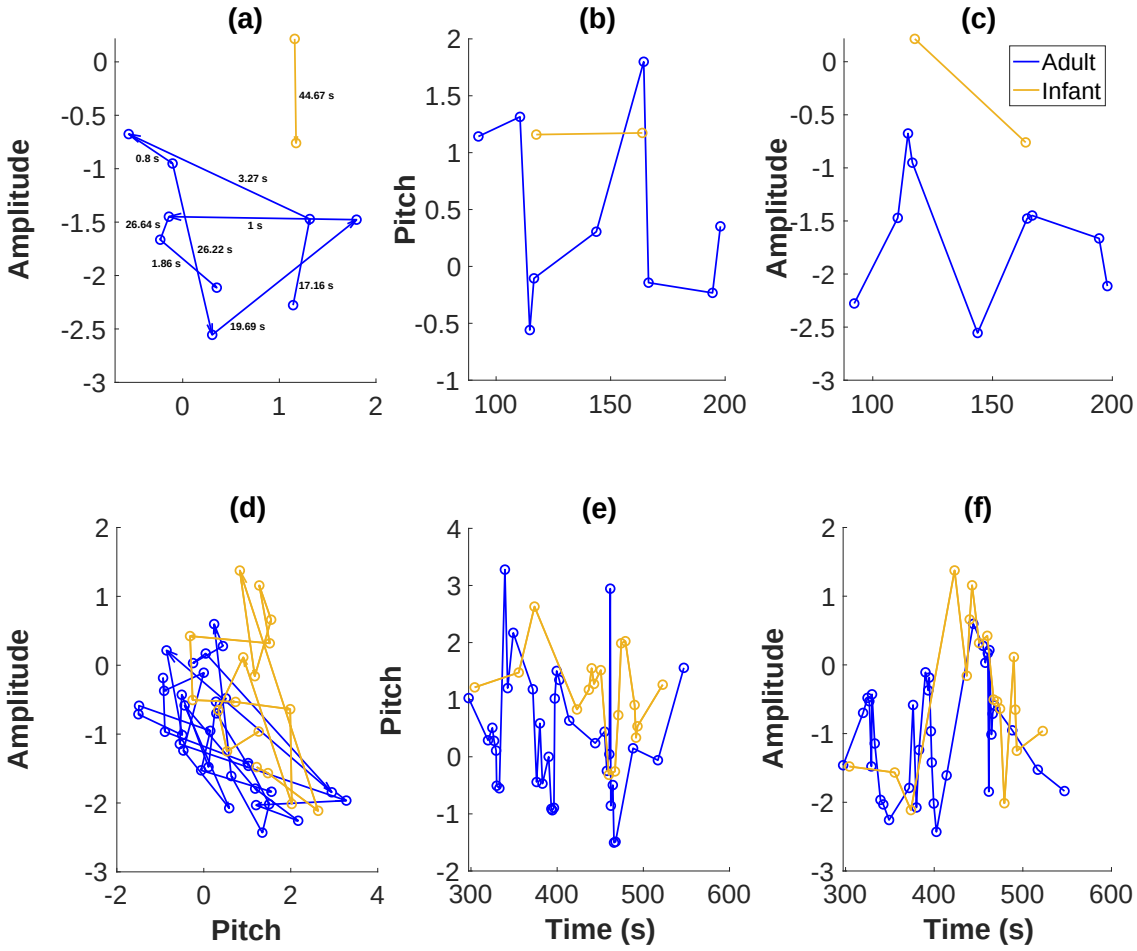

Figure 3.2: **Sample vocalisation sequences by a 3-month-old infant and the infant's caregiver(s).** (a) shows the movement of the infant (yellow) and adult (blue) in a 2D acoustic space defined by mean pitch and amplitude. Directed arrows indicate the direction of the vocaliser's movement in the acoustic space. This can be thought of as a foraging process-like depiction. The inter-vocalisation time intervals are indicated. (b) and (c) show the same data with time plotted on the X axis, and pitch and amplitude respectively on the Y axis. In (a–c), a vocalisation by the adult is indicated by a blue open circle, and a vocalisation by the infant is indicated by a yellow open circle. Note that these are the same data that is depicted in Fig. 3.1. (d), (e), and (f) show plots for a longer period of time taken from the same recording as in (a–c).

Results from a linear mixed effects model with acoustic step size being predicted by inter-vocalisation interval, controlling for age and infant ID, indicate that step sizes in acoustic space and time are positively correlated for both infants ($\beta = .07$, p = < 0.001, n

= 149,272, where β is the standardized regression coefficient) and adults (β = .05, p = < 0.001, n = 460,946). These results also hold for acoustic step sizes and inter-vocalisation intervals following the reception of a response (WR) and following vocalisations that did not receive responses (WOR). For more details, see Table B.4.

At the recording day level, step sizes in 2D acoustic space were found to be predominantly lognormal for both infant and adult vocal exploration for both step types, based on Akaike information criterion (AIC) [113]. Step sizes in individual acoustic dimensions were exponentially distributed for all step types for both speaker types, per AIC (Fig. B.4). Inter-vocalisation intervals between adult vocalisations were predominantly pareto distributed for both step types. A pareto distribution is a type of power law distribution given by $P(x) = \frac{\alpha x_{min}^{\alpha}}{x^{\alpha+1}}$, where $x \in [x_{min}, \infty]$. If the exponent α is such that $1 < \alpha \leq 3$, then a pareto distribution can be used to describe a Lévy walk. Inter-vocalisation times between infant vocalisations were lognormally distributed for both step types. The unsplit step size distributions (i.e., step size distributions that were not split into WR and WOR categories) were largely of the same type as the split step size distributions, except for unsplit time step sizes between adult vocalisations, which were found to be predominantly lognormal (Fig. B.4). For all step size distribution analyses comparing WR to WOR, we only considered steps for which the corresponding inter-vocalisation interval was at least 1 s. This was to control for the fact that determining a response was not received required a wait of 1 s (see Methods).

A randomly selected example of the probability distributions obtained from the day-level recording data compared to the corresponding AIC fit is shown in Fig. B.3. We see reasonable agreement between the two based on visual observation. In addition, we calculated the coefficient of determination (R^2) for all fits computed using AIC and generally found good agreement between raw data and fits. R^2 values indicated especially good fits for WOR and unsplit step size distributions (which had larger sample sizes compared to WR step size distributions), and the mean R^2 value for all fits was found to be 0.73 (Table B.5).

There was no statistically significant interaction between response and age on any of the step size measures. Thus, we do not include an interaction term in any of the statistical tests presented below. For results from statistical tests from models with interaction between response and age effects as a predictor see `https://osf.io/8ern6/`.

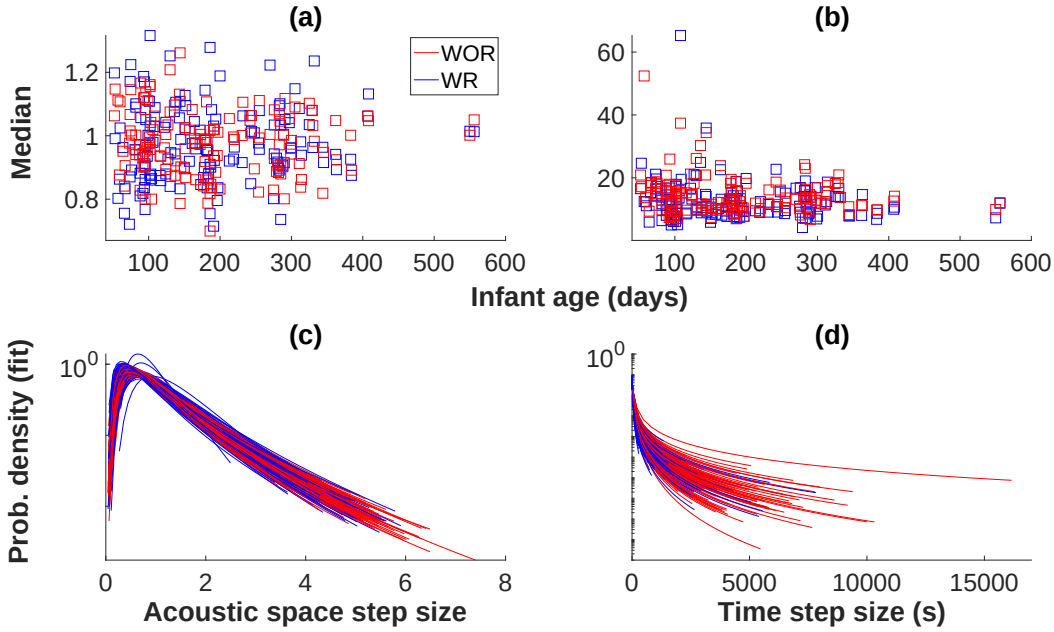

Figure 3.3: **Median values and fitted distributions for acoustic space step size and inter-vocalisation interval distributions for infants.** (a) Median values of infant acoustic step size distributions split into steps following an adult response (WR, blue), and steps following vocalisations that did not receive adult responses (WOR, red). (b) Median values of infant inter-vocalisation interval distributions for WR and WOR steps. Median values were computed based on data, not AIC best fits. (c, d) Infant acoustic step size and inter-vocalisation interval distributions, respectively, fit to lognormal distributions, based on AIC, are shown in log-linear plots. Only distributions that are best fit to lognormal are shown.

3.2.2 Changes with age

Median infant step size in the 2D acoustic space increased significantly with infant age (Table 3.1 and Fig. 3.3; see also Table B.10 and Fig. B.10). Infant step size in the pitch dimension increased as infants got older (Table B.8 and Fig. B.7). However, infant step size in the amplitude dimension if anything got smaller as infants got older (Table B.6 and Fig. B.5). There were no statistically significant effects of age on infant inter-vocalisation time intervals (Table 3.1 and Fig. 3.3; see also Table B.11 and Fig. B.12).

The median as well as the 90^{th} percentile value of the adult vocal step distributions increased with infant age (see Table 3.2) which is in agreement with the increase in the μ of the lognormal fits to adult acoustic step size distributions as infant age increased and with

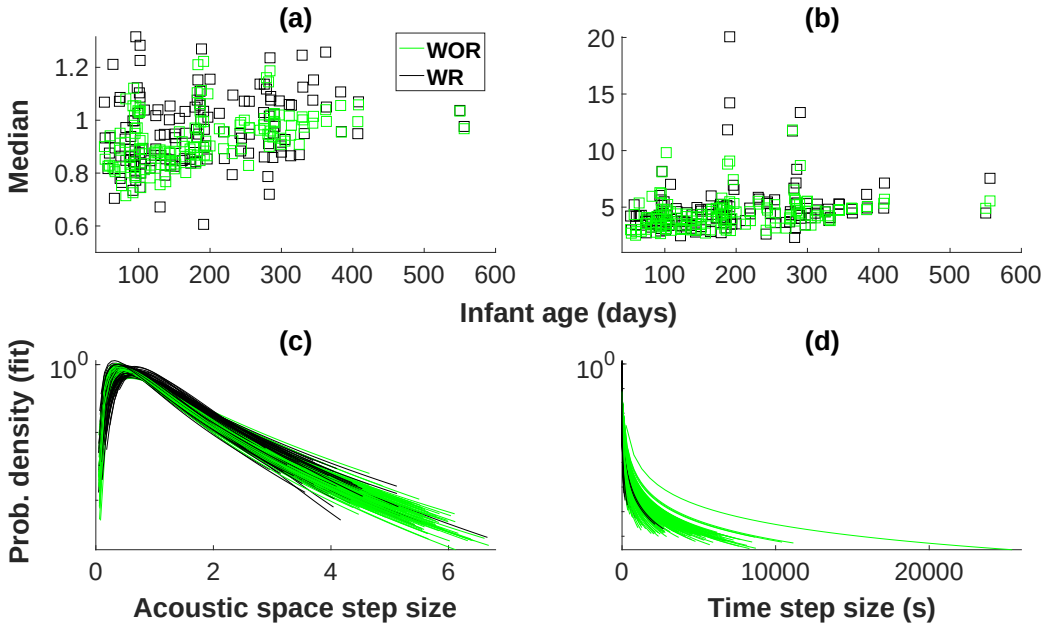

Figure 3.4: **Median values and fitted distributions for acoustic space step size and inter-vocalisation interval distributions for adults.** (a) Median values of adult acoustic step size distributions split into WR (black) and WOR (green). (b) Median values of adult inter-vocalisation interval distributions for WR and WOR. Median values were computed based on data, not AIC best fits. (c, d) Adult acoustic step size and inter-vocalisation interval distributions, respectively. Acoustic step size distributions are fit to lognormal and inter-vocalisation interval distributions are fit to pareto, based on AIC, are shown in log-linear plots. Only distributions that are best fit to lognormal and pareto, respectively are shown.

the decrease in σ (see Fig. B.9a). The changes were evident in both the amplitude and the pitch dimensions (Tables B.7 and B.9). Median inter-vocalisation intervals increased significantly, which is consistent with the decrease in inter-vocalisation interval pareto α values (For a demonstration of how lognormal and pareto distributions change as a function of their parameters, see Fig. B.9 as well as `https://osf.io/2fuje/`). Note that for the analyses of parameters based on step-size distribution fits, we only used those distributions that were best fit to the predominant fit for that distribution type. For example, for μ or σ parameters of adult vocal step distributions, only those distributions that were best fit to lognormal were analysed.

Table 3.1: **Infant steps in 2D acoustic space and inter-vocalisation intervals at the day-long recording level: results of statistical analysis.** Results of mixed effects regressions predicting median and 90th percentile values of vocal foraging measures, as well as parameters of probability distributions of vocal foraging measures, as a function of whether a response was recently received, infant age, and sample size, with infant ID as a random effect. The β values are given with the p values in brackets. Significant effects (at a significance level of 0.05) are indicated in bold. All values reported have been rounded to two decimal points wherever possible.

Dependent variable	**Response effect**	**Age effect**	**Sample size effect**
Inf. acoustic step size (median)	-0.02 (p = 0.85)	**0.15 (p = 0.02)**	**-0.31 (p < 0.001)**
Inf. acoustic step size (90th perc.)	-0.18 (p = 0.13)	0.09 (p = 0.14)	**-0.24 (p = 0.002)**
Inf. acoustic step size (lognormal μ)	0.01 (p = 0.93)	0.12 (p = 0.06)	**-0.30 (p < 0.001)**
Inf. acoustic step size (lognormal σ)	**-0.41 (p < 0.001)**	0.02 (p = 0.70)	**0.18 (p = 0.04)**
Inf. inter-voc. interval (median)	**-0.25 (p = 0.01)**	0.06 (p = 0.22)	**-0.62 (p < 0.001)**
Inf. inter-voc. interval (90th perc.)	-0.11 (p = 0.24)	0.004 (p = 0.94)	**-0.57 (p = 0.01)**
Inf. inter-voc. interval (lognormal μ)	**-0.19 (p = 0.01)**	0.03 (p = 0.52)	**-0.78 (p < 0.001)**
Inf. inter-voc. interval (lognormal σ)	**-0.43 (p < 0.001)**	0.01 (p = 0.76)	**-0.76 (p < 0.001)**

Table 3.2: **Adult steps in 2D acoustic space and inter-vocalisation intervals at the day-long recording level: results of statistical analysis.** Results of mixed effects regressions predicting median and 90$^{\text{th}}$ percentile values of vocal foraging measures, as well as parameters of probability distributions of vocal foraging measures, as a function of whether a response was recently received, infant age, and sample size, with infant ID as a random effect (continued). The β values are given with the p values in brackets. Significant effects (at a significance level of 0.05) are indicated in bold. All values reported have been rounded to two decimal points wherever possible. Note that the p-value indicated by '*' is actually 0.0479 (rounded to 0.05), and therefore, is in bold typeface, indicating significance.

Dependent variable	**Response effect**	**Age effect**	**Sample size effect**
Ad. acoustic step size (median)	**0.24 (p = 0.01)**	**0.39 (p < 0.001)**	0.07 (p = 0.23)
Ad. acoustic step size (90$^{\text{th}}$ perc.)	0.02 (p = 0.83)	**0.29 (p < 0.001)**	**0.19 (p = 0.01)**
Ad. acoustic step size (lognormal μ)	**0.26 (p = 0.003)**	**0.46 (p < 0.001)**	0.05 (p = 0.47)
Ad. acoustic step size (lognormal σ)	**-0.61 (p < 0.001)**	**-0.16 (p = 0.01)**	0.12 (p = 0.11)
Ad. inter-voc. interval (median)	**0.16 (p = 0.05)***	**0.22 (p < 0.001)**	**-0.45 (p < 0.001)**
Ad. inter-voc. interval (90$^{\text{th}}$ perc.)	-0.07 (p = 0.46)	0.04 (p = 0.43)	**-0.50 (p < 0.001)**
Ad. inter-voc. interval (pareto x_{min})	**0.65 (p < 0.001)**	-0.003 (p = 0.96)	-0.04 (p = 0.50)
Ad. inter-voc. interval (pareto α)	**0.30 (p < 0.001)**	**-0.24 (p < 0.001)**	**0.37 (p < 0.001)**

3.2.3 Infant step sizes after receiving an adult response

As shown in Table 3.1 and Fig. 3.3, infants' median and 90$^{\text{th}}$ percentile acoustic space step sizes did not show statistically significant differences as a function of whether the first

vocalisation received a response (see also Table B.10 and Fig. B.10). Upon analysing the parameters of the lognormal fits, we found a smaller σ when a response was received after the first vocalisation in a step. A small σ is associated with a right-shifted peak as well as a smaller tail to the distribution (Fig. B.9a). When looking individually at steps in amplitude (Table B.6, Fig. B.5) and steps in pitch (Table B.8, Fig. B.7), steps in amplitude showed no statistically significant differences as a function of response, and steps in pitch showed only a significantly increased median step size for post-response steps.

Infant inter-vocalisation time intervals were shorter following an adult response, as indicated by statistically smaller median, μ, and σ values (Table 3.1, Fig. 3.3, Table B.11, and Fig. B.12). Infants vocalised again more quickly following the receipt of an adult response.

3.2.4 Adult step sizes after receiving an infant response

For adult vocalisations, the median acoustic step size was significantly larger after an infant response was received compared to other acoustic step sizes, as were μ values (Table 3.2, Fig. 3.4, Table B.10, and Fig. B.10). On the other hand, σ values were significantly smaller for post-response steps. Lower sigma values correspond to distributions with right-shifted peaks but smaller tails. We did not find a significant effect of response for adult 90^{th} percentile 2D acoustic space step size values. Analyses treating amplitude and pitch dimensions individually indicated that steps in the amplitude dimensions increased (larger median and 90^{th} percentile values, and smaller λ values; Table B.7, Fig. B.5) following an infant response whereas steps in the pitch dimension decreased (smaller 90^{th} percentile values and larger λ values; Table B.9, Fig. B.7).

Adult inter-vocalisation intervals tended to decrease following an infant response, as evidenced by larger α values (Table 3.2, Fig. 3.4, Table B.11, and Fig. B.12).

3.2.5 What kinds of vocalisations and vocalisation patterns receive responses?

Table 3.3: **Results of logistic regression predicting whether a response from the other speaker type will follow from a vocalisation's pitch and amplitude.** Infant age was also included as a predictor. Infant ID was included as a random effect. The β values are given with the p values in brackets. Significant effects (at a significance level of 0.05) are indicated in bold. Note that the p-value indicated by '*' is actually 0.0495 (rounded to 0.05), and therefore, is in bold typeface, indicating significance.

Dependent variable	Pitch effect	Amplitude effect	Age effect
Whether infant received adult response	**-0.32** ($p < 0.001$)	**0.17** ($p < 0.001$)	**-0.16** ($p < 0.001$)
Whether adult received infant response	**0.23** ($p < 0.001$)	**0.45** ($p < 0.001$)	**-0.01** ($p = 0.05$)*

As shown in Table 3.3, logistic regression analyses found that infant utterances at lower frequencies and higher amplitudes were more likely to receive adult responses. Overall, the probability of receiving adult responses decreased with infant age. For adults, we found that utterances at higher pitch and amplitude were more likely to receive infant responses, and the overall likelihood of an infant response decreased with infant age, although this effect was only marginally significant (Table 3.3). We also did some preliminary analyses regarding what (non-directional) vocalisation step sizes in pitch, amplitude, and time predict a response. We found the probability of both infants and adults receiving a response increased when a vocalisation was preceded by a larger amplitude step and decreased when a vocalisation was preceded by larger steps in pitch and time (Table B.18).

3.2.6 Validation using data re-labelled by human listeners

Table 3.4: **Inter-rater reliability measures for human-labelled data.** Inter-rater reliability for each human listener with respect to the corresponding LENA labels are presented. All values reported have been rounded to two decimal points wherever possible

Listener id	Child id	Child age	Simple percent agreement	Cohen's Kappa
L1	274	82 days	78.51	0.65
L1	340	183 days	85.61	0.78
L2	530	95 days	71.38	0.53
L3	340	183 days	85.70	0.78

We had four cases where an entire recording's automatically identified child and adult vocalisations were given speaker type labels by listeners who were blind to the speaker type as labelled by the automated system. We calculated inter-rater reliability measures comparing LENA's CHNSP, MAN, and FAN labels to human listener CHN, MAN, and FAN labels, excluding utterances that received any other LENA label or any other human listener label. We found a mean percent agreement of 80.3 (std. dev = 6.83) and a mean Kappa value of 0.68 (std. dev = 0.12) (see Table 3.4). For the data labelled by two human listeners (infant 340, 183 days; listeners L1 and L3), we found a simple percent agreement between those listeners of 0.99 and a Cohen's Kappa value of 0.99, indicating high human inter-rater reliability. For more details on inter-rater reliability measures, see Section B.4.1. We also plotted temporal and acoustic step size distributions for human labelled data compared to automatically labelled data (Fig. B.14–B.17). Visual inspection of these distributions overlaid on each other looked extremely similar, except for CHN WR steps, for which we had very few human-labelled datapoints (see Table B.14). In addition, two sample Kolmogorov-Smirnov tests performed to test the validity of the null hypothesis that step distributions computed from human-labelled data and corresponding LENA-labelled data belong to the same distribution failed to reject the null hypothesis in 78 percent of tests (see Table B.15). We also compared correlation coefficients between step sizes in acoustic space and inter-vocalisation intervals, for data re-labelled by human listeners and the corresponding data as labelled by LENA, and found good agreement between the two (Fig. B.2), with better agreement for the adult data which could be due to there being more adult data points. We

attempted to replicate the results of what vocalisation acoustics predict whether a response will be received, but found that with the human-labelled data, no effects were statistically significant. We hope that future work will attempt to validate the WR vs. WOR and the age results presented here using a larger and more fully annotated human-labelled dataset.

3.3 Discussion

This study set out to examine infant and caregiver vocal behaviour from a foraging perspective. Step size distributions provide one approach to characterising an individual's exploration (and exploitation) processes. They have proved useful in understanding spatial foraging, for example to measure effects of resource scarcity on foraging strategies. It has been shown that adult humans performing non-spatial cognitive tasks, such as coming up with a large list of unique animal names, show movement dynamics in semantic space or on semantic networks that share properties with spatial foraging, such as power law step size distributions and time-distance step size correlations [105].

Inspired by this prior literature on spatial foraging across species and memory foraging by human adults, we asked how infant and caregiver vocalisation-to-vocalisation 'step sizes' change with infant age and depending on whether the speaker's prior vocalisation received a response. Our interest in changes in exploration with age stems from the fact that over the first 18 months of life (the range of ages included in our datasets) infants exhibit enormous changes in the types of vocalisations they produce and in their ability to use vocalisation for a range of communicative purposes. Our interest in whether infant vocal exploration depends on adult responses is based on prior findings that adults respond contingently to infants' non-distress vocalisations, being especially likely to respond vocally to infant vocalisations that are more advanced, and findings that infants' subsequent vocalisations are shaped by recent adult responses [82, 114]. Our interest in whether adult vocal behaviour depends on infant responses is based on prior literature demonstrating that adults actively modify their vocalisation acoustics when they are addressing an infant, and that they do so to entertain and soothe infants as well as to promote language learning. Consistent with prior spatial and memory foraging literature, we characterised distributions of step sizes (a.k.a. lengths) in space, in our case an acoustic space comprised of two dimensions, a vocalisation's mean pitch on log scale and mean intensity, as well as distributions of inter-event intervals, or how much time has passed from the end of one vocalisation to the beginning of the next.

First we tested for correlations between inter-vocalisation time and inter-vocalisation distance in acoustic space. Distance and flight or walk time are naturally correlated in spatial foraging; longer distances take longer to traverse than shorter distances. Prior research on memory foraging has found a correlation between inter-response-intervals and semantic distance [104, 105], and this finding has informed our understanding of the organisation and access of semantic memory. In the present study, we found that correlations between inter-vocalisation intervals and acoustic distance between vocalisations were weak (but very statistically reliable). One possible explanation for the finding that the correlation was weak is that our acoustic space did not provide a great representation of infant and adult vocalisation features. Our 2D acoustic space took into account only the mean pitch and amplitude of each utterance. It ignored all other temporal and spectral features, some of which are key to identifying major types of vocalisations such as canonical (syllablic) infant babble. Errors in labelling the onsets and offsets of vocalisations and in the measurement of pitch could also have played a role in reducing the correlation between inter-vocalisation interval and distance in acoustic space. On the other hand, there is good reason to believe that mean pitch and amplitude are key features that both adults and infants vary in their own vocalisations and attend to in each others' [90,93,94,101]. With more accurate labelling and a more comprehensive acoustic space validated by listener judgments, we may either find that time-space correlations remain weak or that they are revealed to be as strong as those found in memory and spatial foraging. Either way, the results will be highly informative about infant and adult vocal exploration, so such future work will be quite worthwhile.

Inspired by the animal and memory foraging literature, we used AIC to determine whether exponential, normal, lognormal, or pareto distributions best fit the inter-vocalisation step-size distributions for a given recording (and speaker type and response context). We found that step sizes in both time and space were largely lognormally distributed. This suggests that vocal search trajectories are clustered, a feature shared with spatial and memory foraging. On the other hand, the fact that step size distributions were not pareto (the exception being adult inter-vocalisation intervals split according to whether a response was just received) suggests that vocal foraging may not be a Lévy process, as has been found for those other domains. That said, future advancements in vocalisation labelling and measurement will be important to confirm or modify this conclusion. It is also worth noting that we found that when acoustic dimensions were analysed individually, step size distributions were exponentially distributed. This is unsurprising since computing Euclidean distances

on data points distributed exponentially in two orthogonal dimensions produces a probability density that is not exponential. However, the finding highlights that for exploration in more abstract, potentially high-dimensional spaces representing acoustic features (or semantic, visual, or other cognitive features), the distribution type may be less informative than comparisons of how distribution features co-vary with performance or other variables of interest.

Thus, perhaps more important than determining whether or not behaviour generally shows power law scaling is determining whether behaviour dynamics change, qualitatively and/or quantitatively, as a function of other variables of interest, such as infant maturity or resource availability. Inspired by previous studies of foraging behaviour we pursued a distribution fitting and parameter comparison approach to test for differences in vocal foraging dynamics as a function of age and recent resource acquisition, specifically recent receipt of a vocal response. The distribution fits provided a holistic perspective in the sense that they provide a complete characterisation of the shape of the step size or inter-event-interval distributions. On the other hand, we found the parameters of the fits somewhat challenging to interpret, especially in the case of lognormally distributed data, and thus we found that analysing median and 90^{th} percentile values provided a concrete and interpretable way of characterising the distributions. In many cases the two approaches provided converging, complementary evidence for how step size distributions related to response receipt and infant age.

Although we did not have a specific hypothesis regarding how infant age would relate to exploration, we did find intriguing differences in both infant and adult vocal foraging dynamics as a function of age. In particular, we found that as infants got older (and generally more skilled at vocalising) they explored with bigger steps in pitch but smaller steps in amplitude. We also found that with infant age, adult vocalisations to which infants were exposed exhibited longer times between adult vocalisations as well as bigger vocalisation-to-vocalisation differences in the adult pitch and amplitude (and the 2D acoustic space based on these two acoustic features). These findings seem at least in some cases to be related to changes in exploration dynamics rather than simply tracking age-related changes in amplitude and pitch ranges (Table B.2, Table B.3, and Fig. B.1).

A possible explanation for the infant finding is that over the course of the first year and a half of life infants become increasingly skilled at varying the pitch of their vocalisations and produce more active variation along this dimension either for their own interest or

to elicit the interest of their caregivers [71]. A possible explanation for the adult finding is that as infants become more skilled communicators and more cognitively capable and socially aware in general, adults systematically vary their vocalisations in order to stimulate infants [94, 97, 115]. Another possible explanation might be that as infants get older they begin to spend more time physically distant from their caregivers. This could result in the adult vocalisations infants can clearly hear to be somewhat less frequent [116], therefore having longer inter-vocalisation intervals and possibly as a result, correspondingly larger acoustic step sizes. Or it may be that as infants get older, the accuracy of the automatic labelling and measurement increases, revealing greater variation in the acoustic measures. Again, further exploration and development of labelling and acoustic measurement methods will be helpful in clarifying and better understanding these patterns.

As for social effects of receiving a vocal response on vocal foraging dynamics, we hypothesised that vocal responses from adults would function as rewards (among the resources being foraged for) for infants and vice versa. We expected that when a vocalisation was followed by a vocalisation from another speaker type, the speaker's next vocalisation would occur more quickly and would be more similar to the previous vocalisation than when the preceding vocalisation was not followed by a vocal response.

Inter-vocalisation intervals supported foraging for social responses hypothesis: both infants and adults had shorter inter-vocalisation intervals (as indicated by median values and/or fitted step size distribution parameters) when the first vocalisation in a pair was followed by a response than when it was not. This finding fits well with the body of research on naturalistic infant-caregiver vocal interaction based on human-labelled utterance onsets and offsets—it has been found that across cultures infant speech-related vocalisations show bidirectional first order contingencies on each others' occurrence at various ages throughout the first year of life and beyond [99–101]. This agreement between our data and prior work provides additional validation for our new approach focusing on step size distributions. It also provides converging evidence for bidirectional coupling between infant and adult vocalisation processes, and it supports our initial social foraging hypothesis. It should be noted however, that concerns about first order measures' inability to resolve causal pathways (at least when based on non-experimental observations) apply as much to the present study as they do to the prior body of research [101].

As for social effects on step sizes in acoustic space, we did not see a pattern of statistically significant effects that straightforwardly corresponded to our hypothesis in terms of paths

being shorter following a response—the only statistically significant effects of response on infant acoustic path length were a tendency for pitch steps to be *larger* following response and a tendency for steps in 2D space to have right-shifted distribution peaks and lighter tails corresponding to a smaller lognormal σ parameter. Moreover, the effects of infant responses on adult path lengths were also in the opposite direction from what we expected when considering the 2D acoustic space. The finding that adult acoustic step sizes were generally larger following an infant response could reflect the increased variability associated with infant-directed speech. However, since the effect was only apparent in the amplitude dimension and since amplitude is from the infant's perspective and not the adults', the results are perhaps more likely reflecting a tendency for adults to get closer to infants following an infant response. There was one statistically significant effect consistent with our initial hypothesis relating social response to reduced acoustic step size: adult steps in the pitch dimension tended to be shorter following an infant response.

One possible explanation for the null and unexpected results could be that our acoustic space is not comprehensive enough to capture enough of the key features that infants and adults repeat after getting a response [74, 94–96, 117, 118]. Another possibility is that our definition of response receipt was not ideal, not taking into account enough temporal information and not taking into account any features of the responses, such as whether they were positive, who or what the vocalisation was directed toward, what their acoustic properties were, etc. Too much noise in the automatic labelling and the acoustic measurement could also have limited our ability to detect effects. Perhaps longer timescales where behaviour is not viewed at the utterance level but at the level of groups of utterances would also be good to consider (we did do some preliminary area-restricted search inspired plotting that looked at longer series of events, reported in Fig. B.6, B.8, B.11, and B.13). And of course it is possible that infants really are minimally affected by adult responses. We will be more confident drawing conclusions once future work taking into account other possible acoustic spaces and with better human validation for the labelling has been conducted. Nevertheless, we believe the present study provides a useful framework for asking such questions about social effects on vocal exploration dynamics, and our initial results regarding the effects of infant responses on adult pitch and amplitude step sizes are intriguing.

Finally, we found that certain acoustic features of infant and adult vocalisations predicted whether a response would follow. In particular, we found that adult vocalisations that

were higher in amplitude and higher in pitch were more likely to be followed by infant vocalisations. This supports prior literature indicating that acoustically modified infant-directed speech, which tends to have higher pitch together with greater frequency modulation, is more salient and appealing to infants [94, 101, 115, 119]. It also supports our intuition that higher amplitude sounds are more salient to infants. Infant vocalisations were more likely to be followed by an adult vocalisation when they were lower in pitch and higher in amplitude. The higher amplitude finding can perhaps be explained by the greater salience of higher amplitude sounds, although to the extent that greater amplitude is associated with more mature infant vocalisations, that could also be playing a role. The fact that lower pitch infant vocalisations were more likely to be followed by an adult response might also be explained by those vocalisations sounding more adult-like and perhaps more speech-like. Similarly, we found that vocalisations that are further apart in time and pitch are less likely to receive both infant and adult responses, while vocalisations that are further apart in amplitude are more likely to receive both infant and adult responses. One possible explanation for the time step finding is that closely spaced vocalisation events are more noteworthy to the responder than those that are not. It could also be that events are temporally closer when infants and adults are interacting with each other and hence, responses are more likely during these periods of interaction. The findings about step sizes in pitch and amplitude are harder to interpret since these step sizes were non-directional. Repeating these tests with directional step sizes may be more informative. One way to think of larger amplitude steps being more likely to receive a response is that these could correspond to shifts to louder, and hence more noticeable, vocalisations. It will be interesting for future research to explore the underlying reasons for these effects as well as to identify other acoustic predictors of responses and to link these more directly with infant and adult foraging behaviour. For example, it may be that infant vocalisation acoustics which predict adult responses become targets for the infant's vocal exploration even in the absence of an interactive adult caregiver, during times of more solitary vocal play. Our findings on vocalisation acoustics patterns predicting a response are especially interesting when juxtaposed with our findings that adult amplitude steps increase after receiving an infant response, while adult pitch steps decrease following an infant response. For example, it could be that adults repeat similar vocalisation patterns that yielded an infant response in an effort to elicit more responses. On the other hand, a tendency for LENA's automatic labelling software to systematically mislabel infant vocalisations as adult speech and vice

versa could have biased the results on how pitch predicts response in the exact directions observed here. Future validation using a more comprehensively human-labelled dataset is necessary.

3.4 Methods

Audio recordings for this study were obtained using the LENA^{TM} Pro system, which consists of a small single-channel audio recorder and cloth vests each with a chest pocket into which the recorder can be placed. Caregivers were instructed to turn on the recorders at the beginning of the day, place them in the vest pockets, and then put the vests on their infants. The recorders can capture up to 16 hours of audio, including sounds made and heard by the infant.

Recordings included in the present study were obtained as part of two separate data collection efforts. The first involved two infant participants, one learning English and Spanish and the other learning German and English, who were recruited via acquaintance with one of the authors and who began the study at 1–2 months of age and were recorded approximately twice per week until 11–13 months. Recordings were made on days that were convenient for the infants' families, and the infants' parents were instructed to turn on the recorder when the infant woke up in the morning and to turn it off when the infant was put to sleep for the evening. The recorder could be paused as needed for privacy purposes.

The second effort was ongoing at the time of this study and involved 15 participants from the Merced, CA region. Seven children were learning only English, four children were learning both English and Spanish with one of these children having Spanish as the primary language, one was learning English, Spanish, and Sahaptin, one was learning English and German, one was learning English with a small amount of French input, and one was learning English together with another language not specified by the caregiver. Infants were recruited via word of mouth, flyers, and in-person recruitment at a local hospital and community events. The infants in this study were scheduled to be recorded for at least 10 hours during a single day at 3, 6, 9, and 18 months of age. Many of the later recordings had not yet been collected by the time data were prepared for analysis for the present paper. Caregivers were asked to record on days when most of the infant's immediate family members were home and when outings and special events, especially involving a lot of socialising, would be minimal. In most cases, recordings were made on typical weekend

days. Caregivers were instructed to turn on the recorder in the morning, no later than 8 am and to turn off the recorder in the evening, no earlier than 7 pm. Caregivers were told that they could pause the recorders for privacy purposes, but that pause time should not exceed one hour total over the course of the target recording period. Parents were able to have the researchers delete sections of the recording when private events took place and they had not been able to pause the recorder in advance. Caregivers also filled out various questionnaires some of which were to be completed on the day of recording, and were provided in most cases with cash compensation for their time assisting with the study. Since caregivers in both data collection efforts were instructed to complete the recording on a single day, we use the infant's age in days on the day of recording for infant age. However, we have not controlled for possible recordings split over multiple days due to caregiver error, and these recordings, if they exist, have infant age as the infant's age in days on the day of the first subrecording.

Both datasets were collected in accordance with relevant guidelines and regulations. Informed consent was obtained from the infant participants' legal guardians, and data collection protocols were approved by the University of Memphis Institutional Review Board (dataset 1) and by the University of California, Merced Institutional Review Board (dataset 2).

Once recordings were completed, caregivers returned the audio recorders to the researchers, and the audio recordings were uploaded to and processed by the LENA Pro software. The software automatically segments each audio recording into a mutually exclusive and exhaustive set of sound source segments, using a speech recognition algorithm involving a combination of Gaussian mixture and hidden Markov models. The algorithm was previously trained on human-labelled LENA recordings collected by the LENA Research Foundation and is dependent on the infant's age in months. All subsequent analyses focused exclusively on segments labelled "CHN" (clear vocalisations by the infant wearing the recorder [120]) containing at least one "Utt" (a speech-related vocalisation, as opposed to cry, laugh, or vegetative vocalisations), called "CHNSP" in our analyses; "FAN" (clear adult female vocalisations); and "MAN" (clear adult male vocalisations). Note that adult vocalisations collectively refer to FAN and MAN labels, i.e., adult vocalisations are not necessarily from a single adult. A study of the system's sensitivity in detecting various sound sources, conducted by the LENA Research Foundation [120], reported that 82% of portions identified by human transcribers as adults vocalising near the infant were correctly

identified by the system as FAN or MAN, and 2% were mislabelled by the system as infant. For portions identified by human transcribers as infant vocalisations, 76% were correctly labelled as such by the system and 7% were mislabelled as adult. Specificity data were not included in the report.

For three of the day-long recordings, research assistants used custom software (source code available at `https://github.com/tim-shea/ivfcr-server`) to listen to the audio clips that were labelled by the LENA Pro software as either CHN, CXN (near-sounding vocalisations by an infant other than the one wearing the recorder), FAN, or MAN. The research assistants then indicated which speaker types they actually heard during those clips, entering multiple labels for a segment if multiple speaker types were heard. One of the three recordings was labelled by two different people, allowing for comparison of results across raters (one recording, of a 6 month old, was re-labelled by listeners 1 and 3, a recording of a different participant at 3 months was re-labelled by listener 2, and a recording of yet another participant at 3 months was labelled by listener 1). Vocalisations were listened to in the order in which they occurred and were mostly labelled in that same order, except when listeners opted to skip the sounds and return to them later (Note that a small percentage of the segments identified by LENA did not receive human listener labels due to some issues with the software not returning users to vocalisations they had opted to skip and return to later. In all step-size analyses, this would have resulted in a small set of steps that are not between temporally sequential infant-to-infant and adult-to-adult step sizes in the data re-labelled by human listeners). This allowed us to obtain an independent measure of the accuracy of the labels as well as to run analyses on recordings with more accurate labels (human listeners, while more idiosyncratic than the automatic labelling software and not perfectly reliable in their judgments, are generally accepted to be much more accurate than the LENA Pro software) so that results could be compared to LENA-labelled recordings to help identify whether inaccuracies in the automatic labelling could have biased our results. For all validation analyses using data re-labelled by human listeners, we used vocalisations with MAN, FAN, or CHN labels, and without multiple speaker labels only. Further, since all our analyses on LENA-labelled data were carried out on infant vocalisation data with CHNSP labels only, we filtered the CHN labels by human listeners (which comprised of both CHNSP and non-CHNSP vocalisations) using the CHNSP labels by LENA before proceeding with any analyses. MAN and FAN segments that were labelled by humans as CHN were kept and were presumed to be predominantly CHNSP. We are less concerned

with missed vocalisation instances (which due to the nature of the human labelling task were likely quite common) than we are with instances where an adult vocalisation is mistaken for a vocalisation by the infant wearing the recorder and vice versa, as such instances would be more problematic for the research questions posed here.

The automatically obtained vocaliser label data were then used to segment each recording's audio into individual WAV files for each CHN (although only CHNSP were used in our analyses) and AN (FAN or MAN) utterance. A pitch contour for each utterance was then automatically obtained using Praat [121], using the auto-correlation method with a minimum pitch of 75 Hz and a maximum of 1000 Hz for both infant and adult vocalisations. The wide range of possible pitch values was used because it is known that infants and adults interacting with infants tend to vocalise with fundamental frequency that often exceeds the range, on both the low and high ends, of fundamental frequency in typical adult-directed adult speech [117, 118]. Average intensity (a measure of amplitude) was also automatically obtained in Praat for each utterance. In some instances, Praat did not detect a clear pitch for any portion of the utterance; these utterances were subsequently excluded from all analyses. For each utterance, we then obtained the log of the mean pitch (fundamental frequency) in Hz and the mean intensity in dB. Both pitch and amplitude were then converted to z-scores with all infant (CHNSP) and adult (FAN, MAN) vocalisations included in the dataset for standardisation. The standardised log mean pitch and mean intensity were then used to position each utterance in a two-dimensional pitch-amplitude space.

We also used the time stamps of the automatically obtained vocaliser labels to determine whether each infant vocalisation was followed by an adult vocalisation within 1 s following the offset of the infant vocalisation, in which case we (operationally) say that the infant vocalisation received a response. We used the same to determine whether each adult vocalisation was followed within 1 s by an infant speech-related (CHNSP) vocalisation, in which case we say that the adult vocalisation received a response. In cases where two infant vocalisations (CHNSP or non-CHNSP) occurred with less than 1 s separation intervening and no adult vocalisation occupied the intervening time, we marked adult response to the first infant vocalisation as 'not applicable', and the same was done for two adult vocalisations occurring with less than 1 s separating intervening and no infant vocalisation occupying the intervening time [83]. For human re-labelled data, a similar approach but using labels as assigned by human listeners was used to determine whether a vocalisation

received a response or not, or whether a response was 'not applicable' (there were some minor differences in the way 'not applicable' was defined; see analysis code for details). We determined acoustic measures (standardised mean log pitch and standardised amplitude) for human-labelled data by matching start times of vocalisations to LENA-labelled data and using the associated acoustic measures.

Linear mixed effects analyses were run predicting day-level recording step size distribution features. As fixed effects, we entered infant age, the reception of adult/infant response, and sample size (the number of CHNSP, MAN, and FAN vocal events in the recording) into the model. Analyses including an age*response interaction term always returned null effects for the interaction; thus we excluded the interaction term from all analyses reported here. We performed separate analysis with and without the interaction term and the sample size term (see Appendix B). Participant ID was always included as a random effect. For these analyses, any pauses of the recorder were filtered out by removing step sizes corresponding to the step from the end of one recording to the beginning of the next—thus, all data analysed are day-level recordings. All step sizes were also non-directional, i.e., we used absolute values of differences in pitch and amplitude for all analyses.

Finally, we ran logistic mixed effects regressions to determine whether certain vocalisations or vocalisation patterns were more likely to receive responses than others. We used a binary response variable (1 for response received, and 0 for no response) as the dependent variable. To test the relationship between the amplitude and pitch of vocalisations and the probability of a vocalisation receiving a response, we entered infant age, z-scored log pitch of the utterance, and z-scored amplitude of the utterance into the model as fixed effects. To test the relationship between step sizes in pitch, amplitude, and time leading to a vocalisation, and the probability of the vocalisation receiving a response, we added step sizes in z-scored log pitch, z-scored amplitude, and time, from vocalisation $i-1$ to vocalisation i (where vocalisation i is the vocalisation of interest) as fixed effects, in addition to z-scored log pitch of the utterance and z-scored amplitude of the utterance, and infant age. Participant ID was treated as a random effect.

Following pre-processing, MATLAB (R2019a) was used for all analyses except for the mixed effects regressions, for which R (version 3.5.2 – "Eggshell Igloo"), and packages lme4 and lmerTest [122] were used.

Chapter 4

The Fitness Trade-offs of Predation: When to Scavenge and When to Steal

4.1 Introduction

Foraging can be energetically costly, and maintaining energetic reserves is crucial for survival. Consumers must interact with other organisms to obtain energy, and the form of these interactions often depends on trophic level. For example, herbivores such as folivores and frugivores consume plant tissue, but generally do not kill individual plants. Predators, by comparison, must kill their prey prior to energetic replenishment, exposing themselves to substantial risk in the process [123,124]. These risks are multifaceted: prey resources are individually energetically rich, but may be less abundant and patchily distributed in both space and time [124–127]. This may require predators to be on the move and can impose a constant risk of starvation [124]. For larger predators that consume larger prey, the mortality risks associated with prey-handling prior to consumption may be substantial [123]. This is particularly true for larger mammalian carnivores in terrestrial environments because their focal prey tend to have body sizes greater than their own [128, 129], demanding increased risks to access abundant energetic rewards.

Predators are not limited to active hunting but can incorporate a wide array of behavioral strategies that alter the various costs of foraging. These alternative modes of predation

allow the predator to adjust the time and energy it spends acquiring prey, the uncertainty of capturing prey, and the potential for injury when handling prey [123, 129–138]. For example, scavenging involves finding and consuming carrion [139], thereby substituting the time and uncertainties involved in prey handling with those of searching, while lowering the risk of injury [123, 129]. However carrion has less consumable tissue than recently deceased prey, while the tissue that remains may be of lower quality and possibly contaminated by harmful microbiota [140–143]. These qualities of carrion become exaggerated over time, such that the amount of digestible tissue declines via decay while the potential for spoilage increases [144]. Kleptoparasitism, or stealing from a co-occurring predator (which we will refer to as the *competitor*), is another mode of acquisition available to predators [131]. Such behavior eliminates the handling time and mortality risks associated with hunting as well as the diminishing returns of scavenging, yet introduces risks associated with stealing from potentially dangerous competitors [131]. Thus, while the competitor in the kleptoparasitic process may or may not be an individual of the same species [131, 145, 146], such a strategy may be more likely if it poses limited overall risk to the predator [131, 145].

The extent to which these different modes of predation are employed varies across species [139], among individuals within a given species [147], and over the course of an individual's life [148]. The acts of hunting, scavenging, and kleptoparasitism are thus better described as a continuum [139], where the implementation of each mode relative to the others defines a predator's overall strategy. In many cases, the deployment of one mode over another can be opportunistic [130, 139, 149]. For example, great white sharks (*Lamnidae*) actively hunt fur seals [150] while also obtaining significant nutritional gain from scavenging whale carcasses [151]. Albatross (*Diomedeidae*) hunt live squid and fish while supplementing their diet with scavenged prey—often from cetacean vomit [152]. The extent to which different predators rely on scavenged resources can be controversial and both under- or overestimated, depending on taxa [134,153–155]. For example, a standard-bearer of scavenging, the spotted hyena (*Crocuta crocuta*), relies less on scavenged resources in some environments than is widely perceived (up to one-third of diet; [156]), whereas lions (*Panthera leo*) rely more on scavenged resources (up to 35% of diet) than is often assumed [157]. While scavenging is ubiquitous as a strategy of opportunity [139, 158], obligate scavenging is rare because the energetic costs associated with acquiring enough carrion to meet energetic demands is high [134].

Kleptoparasitism as a strategy is largely facultative [131,145,159]. This behavior is more

common among species where predator-prey interactions are easily observed by potential kleptoparasites, and is particularly well-documented among birds [145] and large terrestrial carnivores [137]. For example, skua (*Stercorarius* spp.) are predatory seabirds that are described as specialist kleptoparasites [160], even though resources obtained through active theft remain a minority component of diet [161]. Kleptoparasitism has also been observed in spiders [162], hyenas [137, 163], and marine snails [164], though these behaviors appear to supplement diet, rather than form a primary foraging strategy for most predators. Importantly, the advantages of kelptoparasitism are expected to increase when resources are rare and the length of aggressive contests between predator and competitor are short [165].

Optimizing survival strategies requires balancing the energetic trade-offs and risks associated with different modes of predation. As such, the energetic state of a predator is expected to influence the use of alternative foraging strategies over time. A predator's energetic stores—in the form of body fat for most terrestrial vertebrates—determine the time over which it can survive prior to finding and acquiring additional resources. When a predator has sufficient reserves it may be more likely to seek lower-risk resources that provide less energetic reward. In contrast, predators near starvation may be more likely to engage in riskier foraging behaviors [166, 167]. Since organisms' fat stores (which act as energetic reserves) scale with body size [168], starvation risks are very different for small versus larger predators [129, 169]. Beyond starvation, prey and predator body sizes play a central role in determining predator foraging behavior, influencing energetic expenditures and the mortality risks associated with foraging [170–172]. In fact, the body size ratio of predators and prey is a strong predictor of whether trophic interactions are realized between pairs of species in a community [173–175].

Here, we introduce a stochastic dynamic programming (SDP) framework [176] to examine the conditions resulting in active hunting, scavenging, and kleptoparasitic behaviors for terrestrial carnivores foraging in a stochastic environment over time. We leverage allometric relationships governing energetic costs, physiology, and population densities to evaluate how patterns of predation are likely to change based on the body sizes of the focal predator, prey, and the competitor in kleptoparasitism/scavenging. We then compare our model predictions to behavioral observations for a range of mammalian predators spanning an order of magnitude in body size across several continents.

Our model results point to four key findings that shed light on the nature of alternative predator behavioral strategies. First, we find that whether a predator hunts, scavenges,

or steals is strongly correlated with starvation risk, where higher starvation risk increases the diversity of the predatory modes utilized. Second, the deployment of these modes of predation changes as a function of predator, prey, and competitor body size, where both scavenging and kleptoparasitism become dominant strategies as the size of the prey relative to the predator increases. Third, our model expectations most accurately predict the transition between hunting and alternative modes of predation for a range of terrestrial mammalian predators. Finally, we show that the behavioral transition predicted by the model and captured by field observations follows a scaling law with an exponent near three quarters. Taken as a whole, we suggest that our mechanistic model offers particular insight into the role that behaviors such as scavenging and kleptoparasitism play in the arsenal of strategies available to behaviorally and strategically flexible predators.

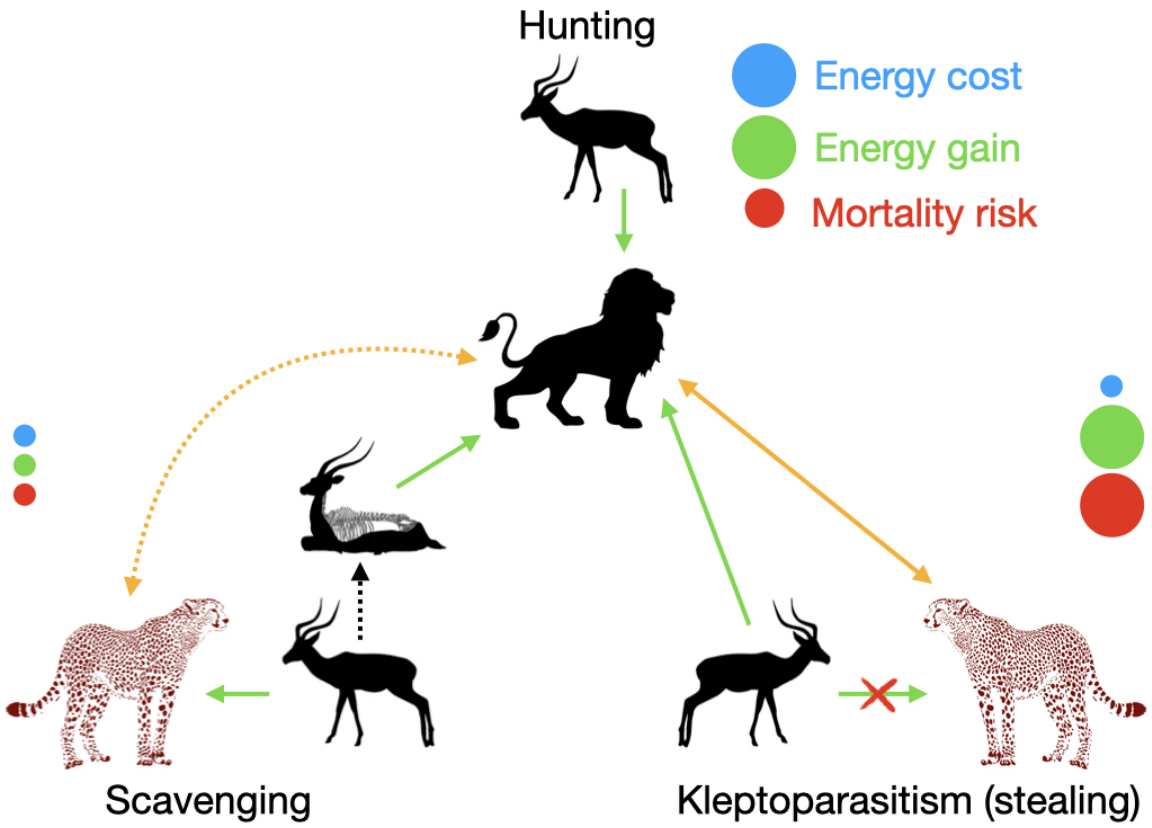

Figure 4.1: **A summary of organismal interactions, energetic transactions, and risks in the model**: The solid orange double-sided arrows indicate direct interactions between predator and kleptoparasitic competitor (brown silhouette) during stealing while the dashed orange double-sided arrow show the indirect interactions between predator and competitor (brown silhouette) during scavenging. Green arrows represent the direction of energy gain. The competitor's interrupted energy gain as a result of stealing by a predator is indicated by the red cross on the energy flow arrow. Finally, the dotted black arrow in the depiction of scavenging represents the transition of freshly-killed prey to leftovers available to the scavenger. The energetic cost, energetic gain, and the mortality risk associated with each predatory mode are represented by blue, green, and red circles, respectively. The size of the circle indicates how big the effect is for the mode of predation in question relative to other modes of predation.

4.1.1 Model Description

We construct a model to explore the conditions under which alternative predator foraging behaviors maximize fitness, where possible behaviors include active hunting, scavenging,

and kleptoparasitism. We assess how these modes of predation change as a function of the risks associated with starvation, mortality induced by intra- or interspecific interactions, and time. To accomplish this, we track two state variables: the energetic state of a focal predator $X = x$ (Joules) and time t (days in a month), where fitness is assessed at the end of the month $t = t_{\max}$. Throughout we maintain uppercase notation for stochastic variables, and lowercase notation for specific values of these stochastic quantities. The focal predator's energetic capacity $X = x_{\max}$, as well as a number of energetic parameters throughout, are based on the body size of the focal predator M_p, its potential prey M_r, and the potential competitor M_c (see Table 4.1). While we introduce the essential framework of the model below, for details please see Section 4.3 and Section C.1.

Hunting, scavenging, and kleptoparasitism have different energetic costs, energetic gains, and mortality risks associated with them (Fig. 4.1). While hunting, a predator searches for the prey resource distributed according to expected densities from Damuth's law [177]. A successful hunting event is an antagonistic encounter (moderate mortality risk, especially for large prey) which involves pursuing and subduing prey before it can be consumed (high energetic cost). Once prey is captured, all of the high quality prey fat and muscle mass is available to the hunter, limited only by its stomach size (high energetic gain).

Kleptoparasitism involves the predator searching for a potential competitor—also distributed according to expected density—and successfully stealing from the competitor if the latter successfully acquires prey through hunting (high energetic gain, same as hunting for the same prey). The low energetic cost due to the lack of the pursue-and-subdue phase is offset by the high mortality risk associated with the antagonistic encounter with another predator. Finally, during scavenging, the focal predator only interacts passively with the competitor and live prey (low energetic cost and mortality risk), but can only consume leftovers from the competitor's meal (low energetic gain).

By virtue of the addition of the competitor, scavenging and stealing are more complex processes than hunting: the success of both behaviors depends on the predator finding a competitor, the competitor successfully hunting prey, and in the case of kleptoparasitism, on the predator finding the competitor before a significant fraction of the prey mass is consumed by the competitor in order to offset the high mortality risk associated with stealing.

Our model incorporates the different ways each strategy utilizes and replenishes predator energetic reserves and associated stochastic risks to relate the predator's updated energetic state to its fitness. All factors described above scale with predator, prey and/or competitor

size, and we use known allometric relationships to parameterize our model as a function of organismal body size.

4.2 Results and Discussion

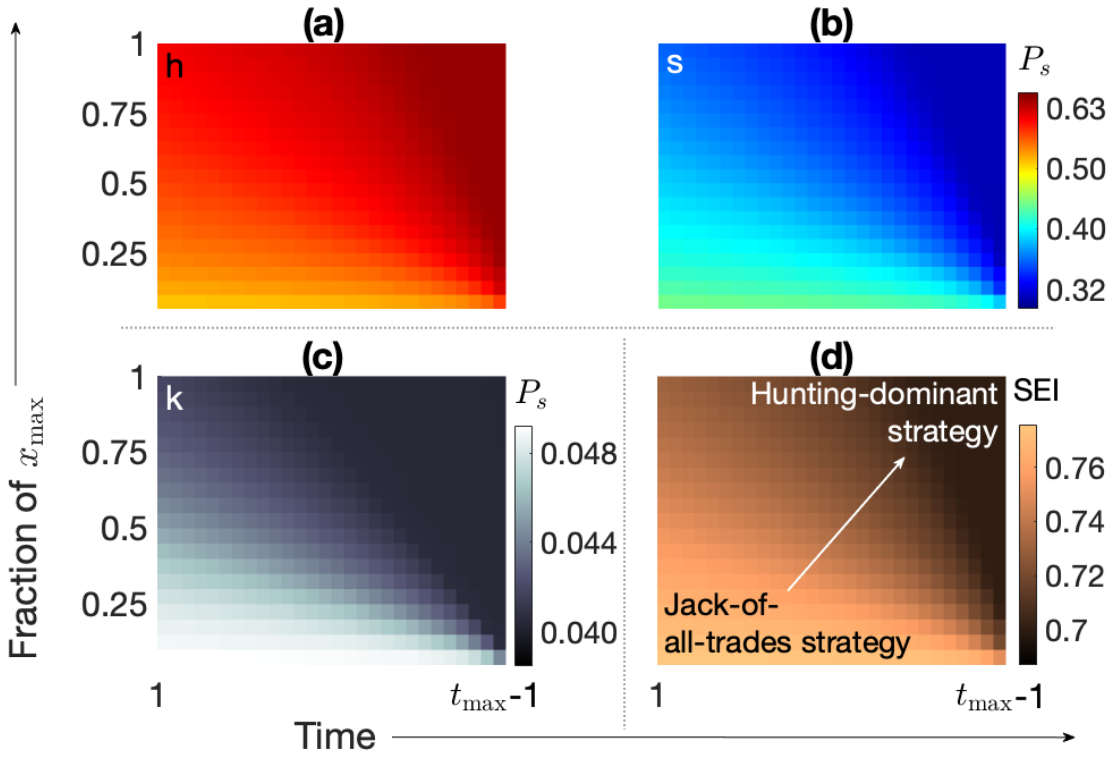

Figure 4.2: **Strategy as a function of predator state.** The average proportion of each strategy (h = hunting, (a); s = scavenging, (b); k = kleptoparasitism, (c)), P_s as a function of the predator's energetic reserves (expressed as a fraction of the maximum energetic storage, $x_{\max}$; X axis) and time (Y axis) are shown. These results are based on decision matrices averaged over all combinations of predator, prey and competitor mass, for 15 trials (see Section C.2 for details). (a) and (b) share a color bar, while (c) has a different color bar, for ease of visualization. Also note that the color maps are on a log scale. Terminal time is denoted by $t_{\max}$, and no decision is made at terminal time. (d) shows the Shannon Evenness Index (SEI; [178–180]) of strategy as a function of the predator's energetic reserves and time.

4.2.1 Predator energetics predict modes of predation

Predator foraging strategies emerge from the cumulative risks associated with finding and successfully acquiring prey while negotiating the potential for serious injury or death. While these behaviors can be viewed through an allometric lens, where the masses of predator, potential prey, and potential competitors scale associated risks, our framework points to important generalities that lay the foundation for this more nuanced perspective. For instance, the results of our model reveal that alternative predator strategies—modes of predation—are strongly predicted by predator energetic state. When the predator's energetic state is replete, there is a heavier reliance on hunting, where reliance is measured by the proportion of states resulting in hunting as the fitness-maximizing strategy averaged across predator, prey, and competitor masses (Fig. 4.2a). As the predator's state declines towards starvation, both scavenging and stealing increase in frequency, where stealing is employed across an order of magnitude fewer states than scavenging (Fig. 4.2b, c). As time advances to the terminal time $t_{\max}$, hunting increases in frequency for a larger proportion of energetic states, with scavenging and to a lesser extent stealing serving as fallback strategies for predators near starvation. These patterns reveal a behavioral switch: near starvation and far from the terminal time, predation modes are more evenly employed, which we denote as the jack-of-all-trades condition (Fig. 4.2d). As the predator attains energetic storage and nears the terminal time, the predator switches to a hunting-dominant condition, which serves to maximize the accessibility of prey while on-boarding modest bodily risk.

The behavioral shift between hunting-dominant to jack-of-all-trades behaviors tracks declining energetic reserves of the predator. A general prediction of our model follows: an increasingly diverse strategic tool-kit is expected to be employed when the risk of starvation-induced mortality is increased. While energetic data on carnivores and their resultant foraging behaviors are very limited, there is some evidence to support such a switch. For example, coyote reliance on ungulate carrion increases during periods when primary prey populations (snowshoe hare) decline [181]. Similarly, Australian dingos [182] and arctic foxes [183] have been observed to rely more on scavenging during periods of resource scarcity, while hyenas [184] have been observed to increase their reliance on scavenging as a result of an effective decrease in prey abundance due to interspecific competition with lions. However, both external environmental and internal physiological drivers of scavenging behavior among predators are not easily quantified [149, 185].

Because predators are more likely to experience near-starvation states in low-productivity

or nutrient-stressed environments, the integration of starvation and/or kleptoparasitic behaviors would increase dependence on—and competition for—carrion subsidies in these conditions. Compellingly, expectations from the stress gradient hypothesis [186], where the negative effects of competition initially increase with environmental stessors, have been proposed to influence predator scavenging behaviors, potentially resulting in higher rates of intraguild competition and significant top-down mesopredator control [185]. In a broad sense, our framework thus offers a mechanistic reasoning for this expectation of carrion subsidy dietary integration that promotes increased intraguild competition. However, it is reasonable to expect that the magnitude of carrion integration should depend on the body sizes of predators, potential prey, and potential competitors, as we will next investigate.

4.2.2 Pairwise allometry constrains modes of predation

Whether a predator hunts, scavenges, or steals is highly constrained by predator, prey, and competitor body sizes. Our model predicts that, on average, larger predators hunt across a larger proportion of states, and maintain this behavior across a larger range of prey body sizes (Fig. 4.3a, d). As expected, as prey increase in size, both scavenging and kleptoparasitic strategies dominate (Fig. 4.3b), whereas we observe a decline in kleptoparasitism as an effective strategy with increasing body size of the competitor (Fig. 4.3c). Smaller to intermediate-sized predators are thus expected to deploy an increased diversity of predator behaviors, while the largest predators tend to hunt. However, these general trends belie the fact that predatory strategies are inherently a function of pairwise interactions between predators and prey, as well as predators and competitors.

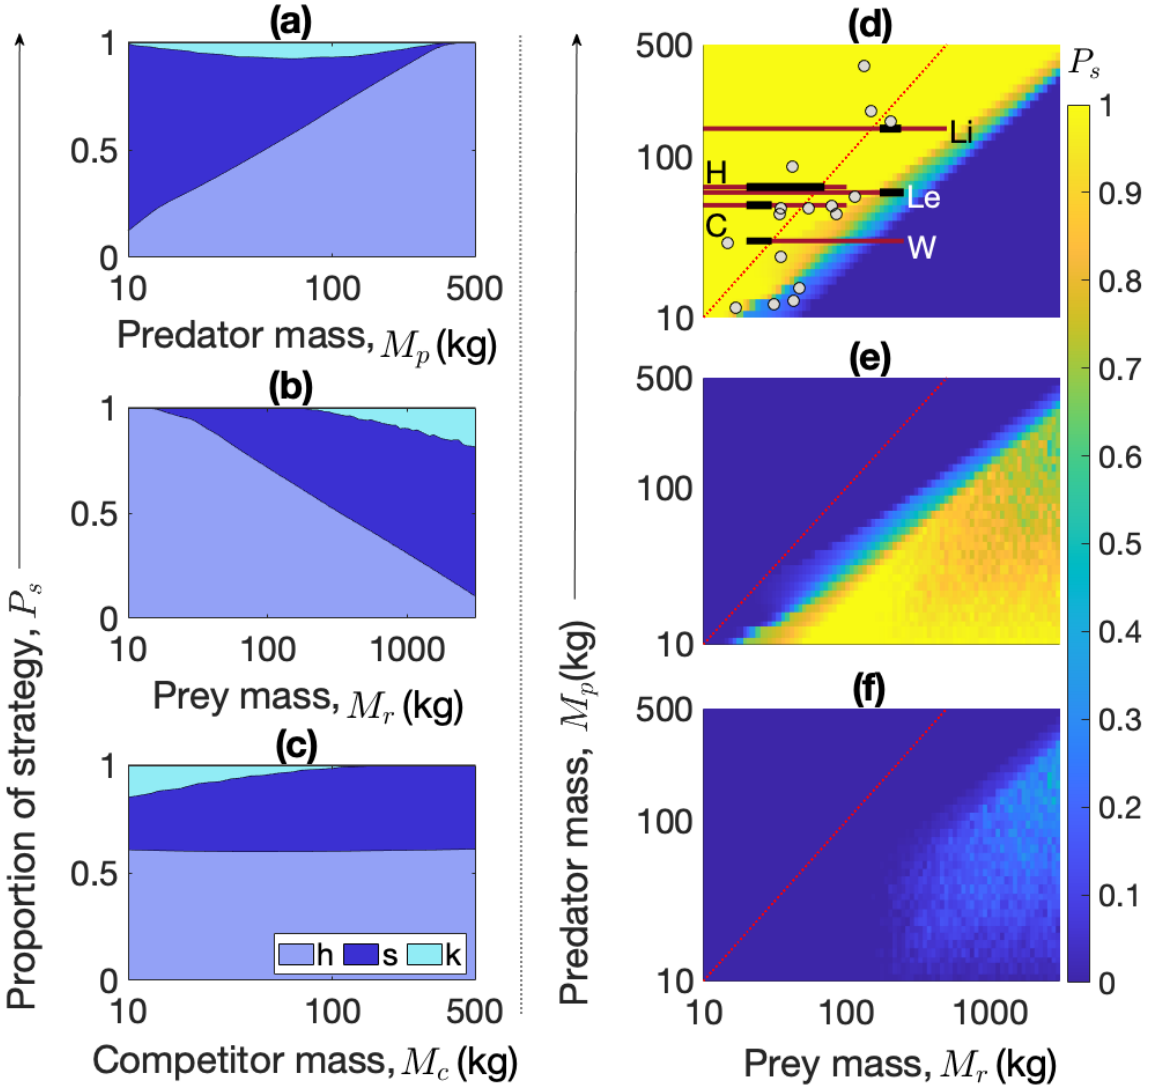

Figure 4.3: **Strategy as a function of body size.** (a)-(c) The proportion of each strategy (h = hunting, s = scavenging, k = kleptoparasitism), P_s as a function of predator mass, M_p, prey mass, M_r, and competitor mass, M_c in kg averaged across 15 replicates (see Section C.2). (d)-(f) Proportion of strategies (P_s, averaged across decision matrices and competitor masses; see Section C.2) resulting in hunting, scavenging, and stealing, respectively, as a function of pairwise predator-prey body size combinations. The dashed line denotes 1:1 predator-prey body sizes. In (d), red and black horizontal bars denote observed total prey and preferred prey mass ranges for predators of different body sizes, respectively [128], where W = African wild dog, C = cheetah, Le = leopard, H = spotted hyena, and Li = lion (Le and H have been staggered for visualization). Grey circles denote the most frequently observed prey for a range of predator body sizes [133].

The fitness advantages of alternative strategies between interacting species reveal allometrically constrained behavioral boundaries. Across predator and prey body sizes, our framework points to a behavioral switch from a hunting-dominant region when the predator is larger than the prey, to a jack-of-all-trades region when the predator is smaller than the prey (Fig. 4.3d-f). This behavioral transition scales sub-linearly with prey body size, meaning that hunting remains the dominant mode of predation for a wider range of prey with increasing predator body size, a trend observed in terrestrial mammalian systems [128]. While our model is relatively coarse and cannot shed light on the nuanced behaviors between species over short timescales, we observe that it successfully predicts predation limitations in diverse mammalian communities. For the most part, observed predator-prey body mass relationships (points and bars in Fig. 4.3d; see Refs. 128, 133) fall within the hunting-dominant strategy space predicted by our framework.

For body size relationships beyond the hunting-dominant region, both scavenging and to a lesser extent kleptoparasitism play increasingly important roles, though observational data for both are limited and constrained to a small number of well-studied species. Kleptoparasitism is employed more frequently for larger prey (Fig. 4.3b), attesting to the increased profitability of stealing larger prey in antagonistic encounters. In contrast, increasing competitor size (proportional to increased mortality risk) results in decreased reliance on kleptoparasitism. Together, we observe that the most diverse strategy tool-kit emerges when prey are roughly 10$\times$ larger than the predator, and when competitors are of similar size or smaller than the predator (Fig. C.10).

Of particular note is the observed nonlinearity in the role of kleptoparasitism and predator body size, with a peak frequency associated with a predator $\approx$ 70 Kg (Fig. 4.3a). This corresponds to a size similar to that of the spotted hyena (*Crocuta crocuta*), a noted kleptoparasite of wild dogs, cheetahs, and lions [187–189]. As an oft-cited exemplar of intraguild instigation, the contributions of hunting, scavenging, and kleptoparasitism to spotted hyena diet are estimated at 50-85%, 7-33% [139], and $\approx$ 20% [187], respectively. Expectations from our model are on par with these observations: for hyena-sized predators, our framework predicts strategy contributions of 59%, 33% and 7%, respectively.

Expanding our assessment of model accuracy against both lions and spotted hyenas, both well-studied large mammalian species engaged in both hunting and scavenging behaviors, we find that model expectations are largely predictive of observed behavioral tendencies. Across sub-Saharan Africa, lion and hyena populations variably supplement active

hunting with scavenging, with both species employing scavenging behaviors from $< 10\%$ to ca. 33% [139]. While lions tend to acquire a greater proportional dietary contribution from activities related to hunting compared to spotted hyenas, the range of the behaviors for both nearly overlap (Fig. 4.4). These field observations of hunting versus scavenging behaviors align with expectations from our model, where we use the proportion of states resulting in a hunting versus scavenging fitness-maximizing strategy—taken across prey body sizes—as a proxy for the percent contribution to diet measured in the field [139]. While our proxy measurement is not one-to-one, we expect it to vary proportionally, such that increases in the proportion of states resulting in scavenging as a fitness maximizing strategy will result in behaviors that tend towards scavenging.

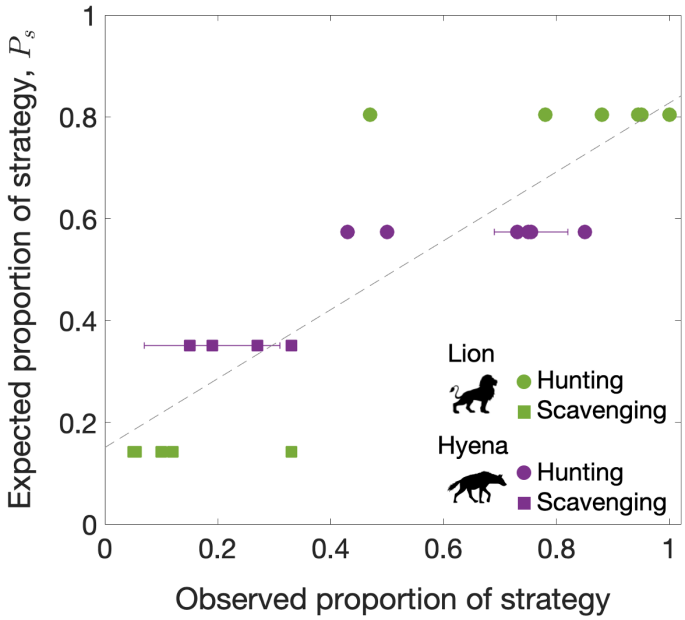

Figure 4.4: **Testing model predictions against data from lion and hyena populations in Africa**. Model predictions for hunting (circles) and scavenging (squares) from lion (green) and spotted hyena (purple) populations across Africa, compiled by [139]. Horizontal error bars indicate maximum and minimum values of the observed strategy proportions where available, as reported in [139], while vertical error bars are too small to be distinguished in the figure. The linear best fit line (dashed gray line) is $y = 0.6769x + 0.1508$ with $R^2 = 0.7953$, indicating good agreement between data and predictions from our model.

4.2.3 The scaling of behavioral transitions

Behaviors are often borne from an intersection of physical and biological constraints, a source from which plastic responses may adapt to rapidly changing situations. Far from these constraints, it is reasonable to expect behaviors to be idiosyncratic and subject to a diversity of (a)biotic drivers. Yet close to the constraint, behaviors may be expected to reflect the nature of the constraint itself. Because our model primarily serves to identify transitions between fitness-maximizing behaviors in a foraging context, we focus assessment of model expectations against observations of these 'behavioral boundaries'. Specifically, we examine the extent to which predators actively hunt prey for a range of predator and prey body sizes, in terms of the percent of kills attributed to prey as well as prey preference (Jacob's index) [190]. While we do not have a direct proxy for these empirical measures in our model, we expect that the transition from hunting to scavenging and/or kleptoparasitism as a function of prey body size will correspond with the prey size where kill percentage and prey preference decline to a minimum. Such an analogy provides an opportunity to examine the predictive capacity of our framework.

Comparing prey kill percentage and preference metrics for six carnivore species spanning nearly an order of magnitude in body size (Fig. 4.5a-f), we observe that our model accurately predicts this behavioral boundary—a transitional prey body size where active hunting by the predator tapers off. The prey body size marking this decline in both kill percentages and prey preference is different for each examined carnivore: it is both predicted and observed to occur at smaller prey size for wild dogs and at much larger prey body size for tigers. We next examine whether our model quantitatively predicts the boundary created by the tapering of observational data that is visually apparent in Fig. 4.5a-f. To extract this boundary for each carnivore species, we use the following procedure.

We first discard prey kill percentages less than 5% because they identify the absence rather than presence of predators' hunting interactions with prey. We then calculate a critical prey body mass M_r^* based on the remaining non-negligible kill percentages. M_r^* is estimated such that all reported prey masses less than or equal to M_r^* cumulatively account for 90% of observations of prey kill percentages. We compare this empirical measure of critical prey mass against the modeled prey mass at which the proportion of hunting falls to 10%. Our expectations of this boundary both visually and quantitatively match observational data for 5 out of 6 species (Fig. 4.5h), as we cannot evaluate the accuracy for wild dogs given the coarseness of kill percentage data available. For this range of species,

the observed boundary reflects an observational limit without shedding light on changes in behavior that might occur for prey body sizes above the threshold, except the obvious decline in active hunting. Our model suggests, however, that at this transition we are more likely to observe alternative modes of predation, a prediction that may be useful in directing future observational efforts.

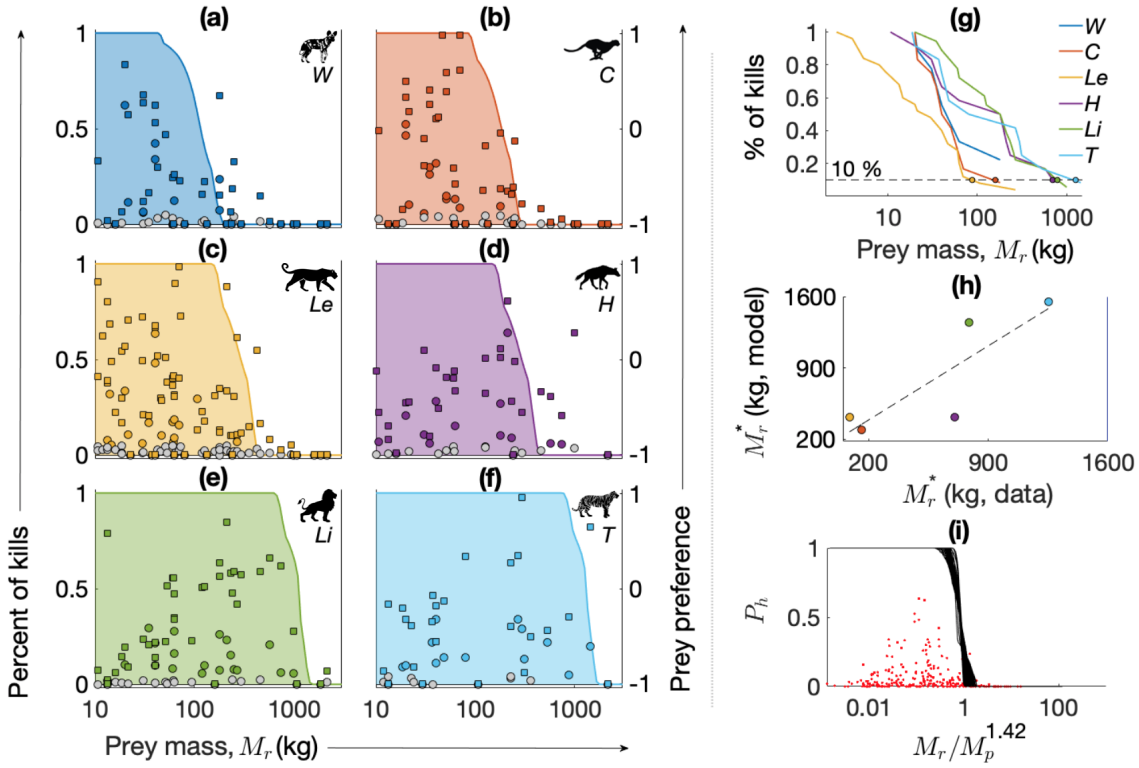

Figure 4.5: **Predicting behavioral transitions in prey preference** (a)-(f) Model predictions of the proportional utilization of hunting (line bordering shaded region) compared against observed prey preferences for (a) wild dog (W), (b) cheetah (C), (c) leopard (Le), (d) spotted hyena (H), (e) lion (Li), and (f) tiger (T). Two empirical measures of prey preference are shown: 1) Jacob's Index of prey (squares), ranging from -1 (avoidance) to +1 (preference) [190] on the right Y axis; and 2) percent contribution of the number of kills of the prey to the total number of kills made by the predator (circles) [191–196] on the left Y axis. Negligible kill percentages (¡ 5%) are indicated by grey circles. (g) Cumulative distribution of non-negligible prey kill percentages shown in (a)-(f) as a function of prey body mass. The 10% threshold is denoted by the gray dashed line while critical prey mass values (M_r^*) estimated from the observational data in (a)-(f) are indicated by circles. (h) Correlation between empirical (from (g); X axis) and predicted (from lines bordering shaded regions in A-F; Y axis) critical prey masses (M_r^*), where hunting dominant behaviors transition to increasingly diverse strategies. The linear best fit line (dashed grey line) is $y = 1.038x + 182.6$, $R^2 = 0.7217$. (i) Rescaled behavioral transition from hunting dominant behaviors to more diverse strategies reveals a common threshold as a function of the prey:predator ratio $M_r/M_p^{1.42}$. Red points denote prey kill percentages for species in A-F.

Perhaps compellingly, while expectations of this behavioral boundary are a function of

the fitness maximization procedure implemented in the SDP (see Section 4.3) and emerge at different prey body sizes for different predators, we observe that they collapse together, such that the behavioral transition is observed to occur at roughly $M_r^*/M_p^{1.42} \approx 1$ (Fig. 4.5i), where M_r^* is the critical prey mass. This points to a scaling relationship for this behavioral boundary where the transitional prey mass $M_r^* \approx M_p^{1.42}$, meaning that predators of larger body size transition from hunting to alternative modes of predation at proportionally larger prey masses. Because larger prey deliver both greater energetic rewards and increased risk to the hunter, we interpret the suspension of hunting behaviors at larger prey body sizes as a signal of the energetic gain outweighing the increased risk. Alternatively we could write $M_p \approx M_r^{*0.704}$. That the exponent is roughly 3/4 suggests that these behavioral boundaries may be a product of established biological scaling relationships.

4.2.4 Conclusion and future directions

The model framework introduced here focuses on the energetics of the focal predator and its interactions with potential prey and competitors. However, we do not take into account habitat-specific parameters such as vegetation cover [197–199] or the effect of group foraging [200,201] (among other things). Even so, our model reproduces key trends from behavioural observations of predator hunting behavior (Fig. 4.5g) and provides a holistic framework to guide future observational efforts of predator foraging behaviors. The success of our model in reproducing observed behavioral patterns, despite the exclusion of group foraging suggests that energetic considerations of the individual predator (in species where group foraging is observed) and its interactions with potential prey and competitors are at least as important in determining foraging strategy as the group structure. Nonetheless, a modified model framework with foraging group size taken into account can provide more insights into the factors that drive the use of alternative foraging strategies. While our model predicts the boundaries of hunting as a function of prey mass remarkably well, it loses some predictive power for prey masses far from this boundary (Fig. 4.5a-f). This is likely because it is designed to predict the fitness maximizing strategy based on interactions between a focal predator-prey-competitor triad rather than the optimal foraging strategy for predators in the presence of an assemblage of potential prey and competitors. However, the model framework can be easily adapted to investigate the adoption of different foraging strategies by predators in real ecological communities without incurring significantly more computational costs, and is the logical next step here.

4.3 Methods

4.3.1 SDP model

We use a simple SDP model to investigate how the energetics of a focal predator and its interactions with potential prey and competitors affect predator foraging behavior. In this model framework, the predator's state is fully specified by its energetic state $X = x$ and time t. In the interest of keeping computational costs manageable, we set the terminal time $t_{\max}$ to 30 days.

To asses predator fitness at $t_{\max}$, we use a saturating terminal fitness function $F(x; t_{\max})$.

$$F(x; t_{\max}) = 1 - e^{x/x_{\max}} \tag{4.1}$$

where $x_{\max}$ is the maximum energetic content of of the focal predator. The saturating dependence on x models the diminishing fitness returns of increasing energetic content close to capacity. The onset of starvation is assumed to occur at $x = 0 \equiv x_c$. Predators do not recover once this critical energetic state, x_c is reached.

Consider a predator whose energetic state at time t is $x(t)$. If the predator's foraging strategy at t is hunting, then

$$x_n^h(t+1) = x(t) - \lambda_n^h + \min(r_n^h, s) \tag{4.2}$$

Here, $x_n^h(t+1)$ is the updated energetic state at time $(t+1)$ given n successful prey encounters by the predator, with corresponding encounter probability p_n^h for $n = 0, 1, 2, .., n_{max}^h$. Successful encounters are those which result in the predator acquiring the prey resource. The maximum number of successful encounters, n_{max}^h, is constrained by s/m_r^h, where m_r^h is the consumable prey mass available to the predator in one successful encounter and s is the predator's stomach size. λ_n^h is the metabolic cost accrued by the predator across n successful encounters, and r_n^h is the prey mass consumed by the predator in n successful encounters.

The predator's projected fitness associated with hunting at time t is given by

$$F^h(t) = \sum_n (1 - \mu_n^h) p_n^h F(x_n^h(t+1), t+1) \tag{4.3}$$

where μ_n^h is the predator mortality associated with hunting for n successful encounters with prey.

Expressions for the updated energetic state and projected fitness for scavenging and kleptoparasitism are analogous to Eqs. 4.2 and 4.3, with superscripts s and k denoting appropriate metabolic costs, mortalities, etc. In the case of kleptoparasitism, a successful encounter is one where the predator is able to steal prey from a competitor in an antagonistic event. In the case of scavenging, however, the predator merely needs to wait for a competitor's leftovers after its encounter with the competitor. Thus, kleptoparasitism involves active predator-competitor interactions, while during scavenging, these interactions are passive. The predator's fitness at time t, then, is given by

$$F(t) = \max\Big(F^h(t), F^s(t), F^k(t)\Big) \tag{4.4}$$

The fitness-maximizing choice at time t determines the predator's foraging strategy given $x(t)$. Since $F^h(t), F^s(t)$, and $F^k(t)$ are all functions of $(t+1)$, the SDP framework allows us to compute the predator's foraging strategy for all x and t given the terminal fitness function.

We use the SDP framework to compute the foraging strategy of (terrestrial mammalian) predators in the 10–500 kg range as a function of x and t, for a range of prey (10–3000 kg) and competitor (10–500 kg) body sizes. We use allometric scaling relationships to incorporate the effects of predator, prey, and competitor body sizes into the model.

The computations are performed independently for each combination of predator, prey and competitor masses and determine the fitness-maximizing strategy for the focal predator within each predator-prey-competitor triad. In order to investigate a large enough number of body size combinations while keeping computational costs low, we used logarithmic increments (base 10) to span the mass ranges reported above. This approach reflects observed trends in nature, since there are fewer predators and prey at higher body masses. We performed 15 trials of the SDP computation for all combinations of predator, prey and competitor masses. All reported results are averages based on these trials.

All computations were performed using MATLAB R2021a and/or R2017b.

Table 4.1: **Summary of allometric scaling relationships used in the SDP model.**

Quantity (Units)	Scaling relationship	Units of body mass	Ref.
Fat mass, f_m (g)	$0.02M^{1.19}$	g	[168, 169, 202]
Muscle mass, m_m (g)	$0.38M$	g	[169, 202]
Skeletal mass, s_m (kg)	$0.061M^{1.09}$	kg	[203]
Resting metabolic rate, λ_r (J/s)	$.018M^{0.75}$	g	[169, 204, 205]
Field metabolic rate, λ_f (J/s)	$0.047M^{0.75}$	g	[169, 205]
Maximal metabolic rate, λ_m (J/s)	$39.597M^{0.872}$	kg	[206]
Body velocity, v (m/s)	$0.33M^{0.21}$	kg	[207]
Reaction distance, d † (m)	$1.62(M_pM_r)^{0.21}$	kg	[207, 208]
Handling time, t_h † (s)	$8912M_p^{-1.02}M_r$	kg	[208]
Stomach size, s (kg)	$0.107M^{1.062}$	kg	[209]
Baseline mortality, μ (s^{-1})	1.2 x $10^{-8}M^{-0.24}$	g	[210]
Population density, ρ (num/km^{-2})	1.15 x $10^4M^{-0.78}$	g	[177]

† Reaction distance and handling time are both dependent on the body sizes of the predator as well as prey, hence the use of M_p and M_r, indicating predator and prey masses, respectively. All other scaling relationships are applicable to both predator and prey, and make use of M, to indicate body mass.

4.3.2 Allometric scaling relationships

For a predator of mass M_p, the fat mass scales as $f_m = 0.02M_p^{1.19}$ [168, 169, 202] and sets the energetic capacity $x_{\max}$ for the predator. In the absence of somatic growth and reproduction, all energetic transactions replenish and diminish only the fat mass. Hence, we use fat mass as a proxy for predator energetic content under these assumptions.

We express predator energetic state as well as all metabolic costs and resource gains in Joules by converting tissue mass to energy and vice-versa. To estimate the average energy yield of bulk tissue, we note that the energy yield of fatty acids is about 38 kJ g^{-1} and that of carbohydrates and proteins is about 17 kJ g^{-1} [211]. For our purposes, we use an intermediate value of 20 kJ g^{-1}. We assume that only 10% of the energy from consumed prey mass is assimilated as predator fat mass, to account for energy lost in biosynthesis [205].

We assume that each 24 hour time step comprises of 12 hours of rest and 12 hours of activity, during which the predator may hunt, scavenge, or steal. The predator's energy expenditure during the resting phase is given by the resting metabolic rate, $\lambda_r = 0.018M_p^{0.75}$ [169,204,205]. For both scavenging and stealing, energy expenditure during the active phase is given by the field metabolic rate, $\lambda_f = 0.047M_p^{0.75}$ [169,205]. For hunting, however, energy expenditure during the active phase depends on the number of successful prey encounters, because the structure of the foraging process is dependent on foraging strategy.

During hunting, the focal predator searches for prey, distributed uniformly according to the latter's equilibrium density [177]

$$\rho_r = 1.15\text{x}10^4 M_r^{0.78} \tag{4.5}$$

The predator's encounter time with prey, t_{enc}, is exponentially distributed.

$$P(t_{\text{enc}}) = (vd\rho_r)e^{-(vd\rho_r)t_{\text{enc}}} \tag{4.6}$$

where $v = 0.33M_p^{0.21}$ is the predator's body velocity [207] and $d = 1.62(M_pM_r)^{0.21}$ is the predator's reaction distance [207, 208] with respect to the prey in question. Predator hunting success follows a binomial distribution characterized by the probability that a food web link exists between the predator and prey as described in [173], with $a = 2.51$, $b = 0.79$, and $g = -0.37$. These values correspond to the fitted body size model for the Serengeti ecosystem [173].

Each successful prey encounter during hunting results in energy expenditure according

to the predator's maximal metabolic rate, $\lambda_m = 39.597 M_p^{0.872}$ [206], for the length of time it takes to pursue and subdue prey prior to consumption. This time is estimated as 20% of prey handling time $t_h = 8912 M_p^{-1.02} M_r$ [208]. Energy expenditure for the remainder of the hunting process, during which the predator does not directly interact with prey, is given by the field metabolic rate, λ_f. Consumable prey mass (which determines energy gain) from a successful hunting encounter, r^h, is given by

$$r^h = f_m + m_m \tag{4.7}$$

where $m_m = 0.38 M_r$ [169, 202] is the prey muscle mass and the superscript h denotes hunting.

During both scavenging and kleptoparasitism, the focal predator searches for a potential competitor, also distributed uniformly according to the latter's equilibrium density. The scaling relationship for the competitor's equilibrium density and the encounter time distribution for predator-competitor encounters are the same as those in Eqs. 4.5 and 4.6, respectively, with the competitor mass M_c replacing prey mass.

While the energetic costs associated with scavenging and kleptoparasitism are the same, the energy gain from successful scavenging and kleptoparasitic encounters are different. The energy gain from a successful instance of kleptoparasitsm, where the focal predator encounters a successful competitor *and* successfully steals from said competitor, is the same as that from a successful hunting encounter (Eq. 4.7). Implicit in this is the assumption that kleptoparasitism is only profitable when the predator is able to usurp all of the high quality consumable prey mass from the competitor. For a successful scavenging encounter, where the predator encounters a successful competitor and consumes its leftovers, the prey mass available for consumption is given by

$$r^s = M_r - (f_m + m_m + s_m) \tag{4.8}$$

where $s_m = 0.061 M_r^{1.09}$ [203] is the prey skeletal mass and the superscript s denotes scavenging. This is based on the assumption that scavengers only have access to leftovers from the competitor's meal.

The maximum energetic gain during a foraging bout is limited by the predator's stomach size $s = 0.107 M_p^{1.062}$ [209], predator energetic state x, and predator energetic capacity $x_{\max}$. Predator stomach size also limits the maximum number of successful encounters,

$n_{\max}$ per foraging bout. Since the energy gain per successful encounter is different for scavenging and hunting/kleptoparasitism, $n_{\max}$ is also different for scavenging and hunting/kleptoparasitism.

To estimate predator mortality associated with different strategies, we note that scavenging does not involve antagonistic interactions with the competitor or prey. Hence, the mortality associated with scavenging, μ^s is given by a baseline mortality $\mu = 1.2\text{x}10^{-8} M_p^{-0.24} \equiv \mu^s$ [210]. By the same reasoning, the mortality associated with hunting and kleptoparasitism should be greater than that for scavenging and should increase with the number of successful encounters with prey and competitor, respectively. Since a number of anatomical, physiological, and behavioural quantities scale with body size, it is reasonable to assume that this increased mortality also scales with the body sizes of the organisms involved. To account for increased mortality during hunting, we add a scaling factor of M_r/M_p to the baseline mortality. Then, the mortality associated with hunting, $\mu^h = (1 + M_r/M_p)\mu$ during the predator's pursuit and subdual of prey, and $\mu^h = \mu$ otherwise. Since kleptoparasitism involves antagonistic interactions with the competitor—which is a predator in its own right—we use a scaling factor of $2M_c/M_p$, such that predator mortality associated with kleptoparasitsim, $\mu^k = (1 + 2M_c/M_p)\mu$ during the predator's interaction with the competitor, and $\mu^k = \mu$ otherwise. The length of time during which this enhanced mortality is applicable for each successful encounter is estimated as 10% of the 'competitor handling time' where we use competitor mass in place of prey mass in the scaling relationship for handling time (Table 4.1).

We compute the fitness associated with each strategy using Eqs. 4.2 and 4.3. To estimate the probability of n successful encounters during a foraging bout for each strategy, p_n^{strategy} for all allowed n, we simulate 50000 independent realisations of the predator's search process for the corresponding strategy.

Table 4.1 provides a summary of all scaling relationships used in this model. For more details about the allometric relationships used in this model, see Section C.1.

Chapter 5

Conclusion

This dissertation presents work that at first glance, might appear disparate, and yet, is linked by a the cohesive thread of looking at complex, emergent foraging phenomena at different scales through the lens of a physicist. These studies are also anchored by their reliance on observational data to bolster the perspectives presented in them and to provide a robust foundation for future investigations.

The minimal one-dimensional model presented in Chapter 2 can be used to analyse observed pattern formation in experimental observations of *Synechocystis* communities as well as simulations based on a two-dimensional biophysical model [39] that has been successful in describing the system. This simple analytical model also offers the possibility to investigate the system beyond the linear approximation we employed. Particularly interesting would be the extension of our model to describe pattern formation in non-linear regimes. Another potential modification would involve reconfiguring the model to account for phototactic movement over longer timescales during which bacterial reproduction becomes relevant.

The framework presented in Chapter 3 by describing the vocal interactions between human infants and their adult caregivers as a foraging process in an acoustic space opens the door to using foraging theory to build a mechanistic understanding of the processes that result in human vocal learning. In particular, insights from this study can be used to build predictive models of infants' vocal foraging process, which in turn may have applications in early identification of speech pathologies and delayed language development, and in understanding cultural differences in human speech by providing an enhanced toolkit to study how babies vocally interact with the world around them.

The computational model presented in Chapter 4 is perhaps the first work of its kind

to provide a predictive framework to understand the underlying reasons for plasticity in predator foraging strategy. In fact, this is merely a first step in understanding how alternative modes of predation affects the dynamics of predator-prey systems and the ecological communities they are embedded in. In addition to extending this model to predict likely modes of predation in realistic assemblages of predator and prey species, this framework can be used to predict optimal group sizes in group foraging. This in turn, can lead the way to a more comprehensive understanding of how alternative modes of predation beyond hunting shape ecological systems over both ecological and evolutionary timescales.

Appendix A

Supplementary Information: A One-Dimensional Model for the Initiation of Fingering Instabilities During Cyanobacterial Phototaxis

A.1 Deriving an analytical expression for the growth rate of instabilities, λ as a function of the wave number, k

From Eq. 2.7, we have

$$\frac{\partial^2 h}{\partial t^2} - \bar{\sigma}\,\frac{\partial^3 h}{\partial t \partial s^2} = -\frac{\partial}{\partial s}\left(\beta\,\frac{\partial h}{\partial s}\,\Gamma^2 - \frac{\alpha}{\bar{f}}\,\Gamma\,\frac{\partial \Gamma}{\partial s}\right) \tag{A.1}$$

where

$$\Gamma = \partial h/\partial t - \bar{\sigma}\,\partial^2 h/\partial s^2 \tag{A.2}$$

To determine the onset of instabilities, we perform linear stability analysis around the uniform solution given by

$$h_p(s,t) = v_0 t + \epsilon \hat{h}(s,t) \tag{A.3}$$

We do this by substituting Eq. A.3 into the expression in Eq. A.1 and retaining terms linear in ϵ. As a first step, we write Eq. A.1 in terms of Γ as a function of h_p.

$$\frac{\partial^2 h_p}{\partial t^2} - \bar{\sigma}\,\frac{\partial^3 h_p}{\partial t \partial s^2} = -\frac{\partial}{\partial s}\Bigg(\underbrace{\beta\,(\Gamma(h_p))^2 \frac{\partial h_p}{\partial s}}_{\text{I}}\;\underbrace{-\frac{\alpha}{\bar{f}}\,\Gamma(h_p)\,\frac{\partial \Gamma(h_p)}{\partial s}}_{\text{II}}\Bigg) \tag{A.4}$$

where $\Gamma(h_p)$ indicates that Γ is a function of h_p.

For simplicity, we will evaluate each term in Eq. A.4 separately. First, let us simplify the left hand side by substituting Eq. A.3 for h_p.

$$\begin{aligned}\frac{\partial^2 h_p}{\partial t^2} - \bar{\sigma}\,\frac{\partial^3 h_p}{\partial t \partial s^2} &= \frac{\partial^2}{\partial t^2}(v_0 t + \epsilon\hat{h}) - \bar{\sigma}\frac{\partial^3}{\partial s^2 \partial t}(v_0 t + \epsilon\hat{h}) \\ &= \epsilon\left(\frac{\partial^2 \hat{h}}{\partial t^2} - \bar{\sigma}\frac{\partial^3 \hat{h}}{\partial s^2 \partial t}\right)\end{aligned} \tag{A.5}$$

Here, $\hat{h}(s,t)$ has been abbreviated as $\hat{h}$. Eq. A.5 is linear in ϵ and hence, there is no need to discard any terms. Evaluating part I of Eq. A.4 using Eqs. A.2 and A.3 gives

$$\begin{aligned}-\frac{\partial}{\partial s}\left(\beta\,(\Gamma(h_p))^2\,\frac{\partial h_p}{\partial s}\right) &= -\beta\frac{\partial}{\partial s}\left(\left(\frac{\partial h_p}{\partial t} - \bar{\sigma}\frac{\partial^2 h_p}{\partial s^2}\right)^2\left(\frac{\partial h_p}{\partial s}\right)\right) \\ &= -\beta\frac{\partial}{\partial s}\left(\left(\frac{\partial}{\partial t}(v_0 t + \epsilon\hat{h}) - \bar{\sigma}\,\frac{\partial^2}{\partial s^2}(v_0 t + \epsilon\hat{h})\right)^2\left(\frac{\partial}{\partial s}(v_0 t + \epsilon\hat{h})\right)\right) \\ &= -\beta\frac{\partial}{\partial s}\left(\left(v_0 + \epsilon\frac{\partial \hat{h}}{\partial t} - \epsilon\bar{\sigma}\frac{\partial^2 \hat{h}}{\partial s^2}\right)^2\left(\epsilon\frac{\partial \hat{h}}{\partial s}\right)\right)\end{aligned}$$

Discarding terms quadratic and above in ϵ, the above expression reduces to

$$\begin{aligned}-\frac{\partial}{\partial s}\left(\beta\,(\Gamma(h_p))^2\,\frac{\partial h_p}{\partial s}\right) &\approx -\beta\frac{\partial}{\partial s}\left(\epsilon v_0^2\frac{\partial \hat{h}}{\partial s}\right) \\ &= -\epsilon\beta v_0^2\frac{\partial^2 \hat{h}}{\partial s^2}\end{aligned} \tag{A.6}$$

Part II of Eq. A.4 can be simplified similarly.

$$\begin{aligned}
\frac{\partial}{\partial s}\left(\frac{\alpha}{\bar{f}}\,\Gamma(h_p)\,\frac{\partial\Gamma(h_p)}{\partial s}\right) &= \frac{\partial}{\partial s}\left(\frac{\alpha}{\bar{f}}\left(\frac{\partial h_p}{\partial t}-\bar{\sigma}\frac{\partial^2 h_p}{\partial s^2}\right)\left(\frac{\partial}{\partial s}\left(\frac{\partial h_p}{\partial t}-\bar{\sigma}\frac{\partial^2 h_p}{\partial s^2}\right)\right)\right)\\
&= \frac{\partial}{\partial s}\left(\frac{\alpha}{\bar{f}}\left(\frac{\partial}{\partial t}(v_0 t+\epsilon\hat{h})-\bar{\sigma}\frac{\partial^2}{\partial s^2}(v_0 t+\epsilon\hat{h})\right)\right.\\
&\qquad\left.\left(\frac{\partial^2}{\partial s\partial t}(v_0 t+\epsilon\hat{h})-\bar{\sigma}\frac{\partial^2}{\partial s^2}(v_0 t+\epsilon\hat{h})\right)\right)\\
&= \frac{\partial}{\partial s}\left(\frac{\alpha}{\bar{f}}\left(v_0+\epsilon\frac{\partial\hat{h}}{\partial t}-\epsilon\bar{\sigma}\frac{\partial^2\hat{h}}{\partial s^2}\right)\left(\frac{\partial}{\partial s}\left(v_0+\epsilon\frac{\partial\hat{h}}{\partial t}-\epsilon\bar{\sigma}\frac{\partial^2\hat{h}}{\partial s^2}\right)\right)\right)\\
&= \frac{\partial}{\partial s}\left(\frac{\alpha}{\bar{f}}\left(v_0+\epsilon\frac{\partial\hat{h}}{\partial t}-\epsilon\bar{\sigma}\frac{\partial^2\hat{h}}{\partial s^2}\right)\left(\epsilon\frac{\partial^2\hat{h}}{\partial t\partial s}-\epsilon\bar{\sigma}\frac{\partial^3\hat{h}}{\partial s^3}\right)\right)
\end{aligned}$$

With only terms linear in ϵ retained, the expression above becomes

$$\begin{aligned}
\frac{\partial}{\partial s}\left(\frac{\alpha}{\bar{f}}\,\Gamma(h_p)\,\frac{\partial\Gamma(h_p)}{\partial s}\right) &\approx \epsilon\frac{\alpha v_0}{\bar{f}}\frac{\partial}{\partial s}\left(\frac{\partial^2\hat{h}}{\partial t\partial s}-\bar{\sigma}\frac{\partial^3\hat{h}}{\partial s^3}\right)\\
&= \epsilon\frac{\alpha v_0}{\bar{f}}\frac{\partial^2}{\partial s^2}\left(\frac{\partial\hat{h}}{\partial t}-\bar{\sigma}\frac{\partial^2\hat{h}}{\partial s^2}\right)\equiv\epsilon\frac{\alpha v_0}{\bar{f}}\frac{\partial^2\hat{\Gamma}}{\partial s^2}
\end{aligned}\tag{A.7}$$

where $\hat{\Gamma}=\partial\hat{h}/\partial t-\bar{\sigma}\,\partial^2\hat{h}/\partial s^2$. Rewriting Eq. A.4 in terms of Eqs. A.5–A.7 and dividing by ϵ yields

$$\frac{\partial^2\hat{h}}{\partial t^2}-\bar{\sigma}\frac{\partial^3\hat{h}}{\partial t\partial s^2}=\frac{-\partial^2}{\partial s^2}\left(\beta v_0^2\,\hat{h}-\frac{\alpha v_0}{\bar{f}}\left(\frac{\partial\hat{h}}{\partial t}-\bar{\sigma}\frac{\partial^2\hat{h}}{\partial s^2}\right)\right)\tag{A.8}$$

Substituting solutions of the form $\hat{h}\sim e^{\lambda t}e^{iks}$ into Eq. A.8, we get

$$\begin{aligned}
&\lambda^2\hat{h}-\bar{\sigma}(ik)^2\lambda\hat{h}=-\beta v_0^2(ik)^2\hat{h}+\frac{\alpha v_0}{\bar{f}}\frac{\partial^2}{\partial s^2}(\lambda\hat{h}-\bar{\sigma}(ik)^2\hat{h})\\
\implies\ &\lambda^2\hat{h}+\bar{\sigma}k^2\lambda\hat{h}=\beta v_0^2k^2\hat{h}+\frac{\alpha v_0}{\bar{f}}\big((ik)^2\lambda\hat{h}-\bar{\sigma}(ik)^4\hat{h}\big)\\
\implies\ &\lambda^2+\bar{\sigma}k^2\lambda=\beta v_0^2k^2-\frac{\alpha v_0}{\bar{f}}(k^2\lambda+\bar{\sigma}k^4)
\end{aligned}$$

Here, we have divided the entire equation by $\hat{h}$ in the last line. Rewriting the above expression in the standard quadratic form gives

$$\lambda^2 + \left(\frac{\alpha v_0}{\bar{f}} + \bar{\sigma}\right)k^2\lambda + \frac{\alpha v_0}{\bar{f}}\bar{\sigma}k^4 - \beta v_0^2 k^2 = 0$$

Using $\bar{f} = f\,\alpha/\beta$ and $v_0 = \bar{f}c_0$, this can be rewritten as

$$\lambda^2 + \big(D_0 + \bar{\sigma}\big)k^2\lambda + D_0\bar{\sigma}k^4 - \frac{f^2 D_0^2}{\beta}k^2 = 0 \tag{A.9}$$

where $D_0 = \alpha\, c_0$ and $\beta\, v_0^2 = f^2\, D_0^2/\beta$. λ can be expressed as a function of k by solving the quadratic equation above.

$$\lambda = \frac{-k^2(D_0 + \bar{\sigma})}{2} \pm \sqrt{(D_0 + \bar{\sigma})^2 k^4 - 4\left(D_0\bar{\sigma}k^4 - \frac{f^2 D_0^2 k^2}{\beta}\right)} \tag{A.10}$$

Setting $k_c^2 = f^2 D_0/(\bar{\sigma}\beta)$, we can rewrite Eq. A.10 as

$$\begin{aligned}
\lambda &= \frac{-k^2(D_0 + \bar{\sigma})}{2} \pm \sqrt{\big(D_0 + \bar{\sigma}\big)^2 k^4 - 4\big(D_0\bar{\sigma}k^4 - k_c^2 D_0\bar{\sigma}k^2\big)} \\
&= \frac{-k^2(D_0 + \bar{\sigma})}{2} \pm \sqrt{\big(D_0 + \bar{\sigma}\big)^2 k^4 + 4k^4 D_0\bar{\sigma}\left(\frac{k_c^2}{k^2} - 1\right)} \\
&= \frac{-k^2(D_0 + \bar{\sigma})}{2} \pm \sqrt{\big(D_0 + \bar{\sigma}\big)^2 k^4\left(1 + \frac{4D_0\bar{\sigma}}{(D_0 + \bar{\sigma})^2}\left(\frac{k_c^2}{k^2} - 1\right)\right)}
\end{aligned}$$

This expression reduces to Eq. 2.11 by using $\gamma = 4D_0\bar{\sigma}\big(D_0 + \bar{\sigma}\big)^{-2}$.

$$\lambda = \frac{k^2(D_0 + \bar{\sigma})}{2}\left[-1 \pm \sqrt{1 + \gamma\left(\frac{k_c^2}{k^2} - 1\right)}\right] \tag{A.11}$$

A.2 Analytical expressions for the fastest growing mode, $k_{\max}$ and its growth rate, $\lambda_{\max}$

To find the fastest growing mode, we set $\partial\lambda/\partial k = 0$.

$$\frac{\partial\lambda}{\partial k} = \frac{\partial}{\partial k}\left(\frac{k^2(D_0+\bar{\sigma})}{2}\left[-1 \pm \sqrt{1+\gamma\left(\frac{k_c^2}{k^2}-1\right)}\right]\right) = 0 \tag{A.12}$$

For simplicity, let $K = k_c^2/k^2$, and $A = 1+\gamma(k_c^2/k^2-1) = 1+\gamma(K-1)$. Then, Eq. A.12 becomes

$$\frac{\partial\lambda}{\partial k} = k(D_0+\bar{\sigma})(-1\pm\sqrt{A}) \pm \frac{k^2(D_0+\bar{\sigma})}{2}\left(\frac{1}{2\sqrt{A}}\frac{\partial A}{\partial k}\right) = 0 \tag{A.13}$$

Using $\partial A/\partial k = -2\gamma k_c^2/k^3$ in A.13

$$\begin{aligned}
0 &= (D_0+\bar{\sigma})\left(k(-1\pm\sqrt{A}) \mp \frac{k^2}{2}\left(\frac{2\gamma}{2\sqrt{A}}\frac{k_c^2}{k^3}\right)\right) \\
&= (D_0+\bar{\sigma})\left(k(-1\pm\sqrt{A}) \mp \frac{1}{2}\left(\frac{\gamma}{\sqrt{A}}\frac{k_c^2}{k}\right)\right) \\
&= (D_0+\bar{\sigma})\frac{k_c^2}{k}\left(\frac{k^2}{k_c^2}(-1\pm\sqrt{A}) \mp \frac{\gamma}{2\sqrt{A}}\right)
\end{aligned}$$

Dividing by $(D_0 + \bar{\sigma})k_c^2/k$ throughout and using $K = k_c^2/k^2$ gives

$$\begin{aligned}
0 &= \frac{-1 \pm \sqrt{A}}{K} \mp \frac{\gamma}{2\sqrt{A}} \\
\implies 0 &= \frac{-\sqrt{A} \pm A}{K} \mp \frac{\gamma}{2} \\
\implies \frac{\sqrt{A}}{K} &= \pm\frac{A}{K} \mp \frac{\gamma}{2} \\
\implies \frac{A}{K^2} &= \left(\frac{A}{K} - \frac{\gamma}{2}\right)^2 \\
\implies A &= \left(A - \frac{\gamma K}{2}\right)^2 \\
\implies A &= A^2 + \frac{\gamma^2 K^2}{4} - A\gamma K
\end{aligned}$$

Substituting for A, we get

$$\begin{aligned}
1 + \gamma(K-1) &= 1 + \gamma^2(K-1)^2 + 2\gamma(K-1) + \frac{\gamma^2 K^2}{4} - (1 + \gamma(K-1))\gamma K \\
\implies \gamma(K-1) &= \gamma^2(K-1)^2 + 2\gamma(K-1) + \frac{\gamma^2 K^2}{4} - (1 + \gamma(K-1))\gamma K \\
\implies 0 &= \gamma^2(K-1)^2 + \gamma(K-1) + \frac{\gamma^2 K^2}{4} - (1 + \gamma(K-1))\gamma K \\
&= (\gamma^2 K^2 + \gamma^2 - 2\gamma^2 K) + (\gamma K - \gamma) + \frac{\gamma^2 K^2}{4} + (\gamma^2 K - \gamma K - \gamma^2 K^2)
\end{aligned}$$

Rearranging the above expression yields a quadratic equation in K.

$$\frac{\gamma^2 K^2}{4} - \gamma^2 K + \gamma(\gamma - 1) = 0 \tag{A.14}$$

Since this is obtained by rearranging the equation for $\partial\lambda/\partial k = 0$, the value of k that satisfies Eq. A.14 will also be a solution for $\partial\lambda/\partial k = 0$.

$$\begin{aligned} K &= \frac{\gamma^2 \pm \sqrt{\gamma^4 - \gamma^3(\gamma-1)}}{(\gamma^2/2)} \\ &= \frac{\gamma^2 \pm \sqrt{\gamma^4\left(1 - \frac{(\gamma-1)}{\gamma}\right)}}{(\gamma^2/2)} \\ &= \gamma^2\left(\frac{1 \pm \sqrt{1 - 1 + \frac{1}{\gamma}}}{(\gamma^2/2)}\right) \\ \implies K &= 2\left(1 \pm \frac{1}{\sqrt{\gamma}}\right) \end{aligned} \tag{A.15}$$

Using $K = k_c^2/k^2$, we get

$$k = \frac{k_c}{\sqrt{K}} \tag{A.16}$$

Substituting Eq. A.16 in the positive solution for K in Eq. A.15, we get

$$k_{\max} = k_c\left(2\big(1 + 1/\sqrt{\gamma}\big)\right)^{-\frac{1}{2}} \tag{A.17}$$

Using $k_c^2 = f^2 D_0/(\bar{\sigma}\beta)$ and $\gamma = 4D_0\bar{\sigma}/(D_0 + \bar{\sigma})^2$, we can rewrite Eq. A.17 as

$$\begin{aligned} k_{\max}^2 &= \frac{f^2 D_0}{\bar{\sigma}\beta}\frac{\sqrt{4D_0\bar{\sigma}}}{(D_0 + \bar{\sigma})} \div 2\left(\frac{\sqrt{4D_0\bar{\sigma}}}{(D_0 + \bar{\sigma})} + 1\right) \\ &= \frac{f^2 D_0^{\frac{3}{2}}}{\sqrt{\bar{\sigma}}\beta} \div \left(\sqrt{4D_0\bar{\sigma}} + D_0 + \bar{\sigma}\right) \\ &= \frac{f^2 D_0^{\frac{3}{2}}}{\sqrt{\bar{\sigma}}\beta} \div \left(\sqrt{D_0} + \sqrt{\bar{\sigma}}\right)^2 \\ &= \frac{f^2 D_0^{\frac{3}{2}}}{\sqrt{\bar{\sigma}}\beta\left(\sqrt{D_0} + \sqrt{\bar{\sigma}}\right)^2} \end{aligned}$$

Taking the square root, we get

$$k_{\max} = \frac{f D_0^{\frac{3}{4}} \bar{\sigma}^{-\frac{1}{4}}}{\sqrt{\beta}\left(\sqrt{D_0} + \sqrt{\bar{\sigma}}\right)} \tag{A.18}$$

which is the same expression as in Eq. 2.13. Next, we will derive Eq. 2.14 for the growth rate of the fastest growing mode, $\lambda_{\max}$. To do this, we set $k = k_{\max}$ in Eq. A.11.

$$\lambda_{\max} = \frac{k_{\max}^2 (D_0 + \bar{\sigma})}{2}\left[-1 \pm \sqrt{1 + \gamma\left(\frac{k_c^2}{k_{\max}^2} - 1\right)}\right] \tag{A.19}$$

First, we evaluate the expression inside the square root using Eq. A.17.

$$\begin{aligned}
1 + \gamma\left(\frac{k_c^2}{k_{\max}^2} - 1\right) &= 1 + \gamma\left(\frac{k_c^2}{k_c^2 (2(1 + 1/\sqrt{\gamma}))^{-1}} - 1\right) \\
&= 1 + \gamma\left(2 + \frac{2}{\sqrt{\gamma}} - 1\right) \\
&= 1 + \gamma\left(1 + \frac{2}{\sqrt{\gamma}}\right) \\
&= 1 + \gamma + 2\sqrt{\gamma}
\end{aligned}$$

Substituting this expression in Eq.A.19

$$\begin{aligned}
\lambda_{\max} &= k_{\max}^2 \frac{(D_0 + \bar{\sigma})}{2}\left(-1 \pm \sqrt{1 + \gamma + 2\sqrt{\gamma}}\right) \\
&= k_{\max}^2 \frac{(D_0 + \bar{\sigma})}{2}\left(-1 \pm \sqrt{(1 + \sqrt{\gamma})^2}\right) \\
&= k_{\max}^2 \frac{(D_0 + \bar{\sigma})}{2}\left(-1 \pm (1 + \sqrt{\gamma})\right)
\end{aligned}$$

After discarding the negative solution, we are left with

$$\begin{aligned}
\lambda_{\max} &= k_{\max}^2 \frac{(D_0 + \bar{\sigma})}{2}\left(-1 + (1 + \sqrt{\gamma})\right) \\
&= k_{\max}^2 \frac{(D_0 + \bar{\sigma})}{2}\sqrt{\gamma})
\end{aligned}$$

Using Eq. A.18 and $\gamma = 4D_0\bar{\sigma}\left(D_0+\bar{\sigma}\right)^{-2}$, this becomes

$$\begin{aligned}\lambda_{\max} &= \left(\frac{f^2 D_0^{\frac{3}{2}}\bar{\sigma}^{-\frac{1}{2}}}{\beta\left(\sqrt{D_0}+\sqrt{\bar{\sigma}}\right)^2}\right)\left(\frac{D_0+\bar{\sigma}}{2}\right)\left(\frac{4D_0\bar{\sigma}}{\left(D_0+\bar{\sigma}\right)^2}\right)\\ &= \frac{f^2 D_0^{\frac{3}{2}}\sqrt{D_0\bar{\sigma}}}{\sqrt{\bar{\sigma}}\beta\left(\sqrt{D_0}+\sqrt{\bar{\sigma}}\right)^2}\end{aligned}$$

From this, we get

$$\lambda_{\max} = \frac{f^2 D_0^2}{\beta\left(\sqrt{D_0}+\sqrt{\bar{\sigma}}\right)^2} \tag{A.20}$$

Appendix B

Supplementary Information: Exploratory Dynamics of Vocal Foraging During Infant-Caregiver Communication

Table B.1: **List of abbreviations in figure legends and axis labels**

Abbreviation	Explanation
Ch	Infant vocalisation
Ad	Adult vocalisation
WR	With response (for steps following vocalisations that received responses)
WOR	Without response (for steps following vocalisations that did not receive responses)
LENA or L	Labelled by LENA software
HUM or H	Labelled by human listeners
(f)	AIC best fit curve
(d)	Raw data-based curve
s.z.	Step size
std. dev.	Standard Deviation

B.1 Infant and adult acoustics as a function of infant age

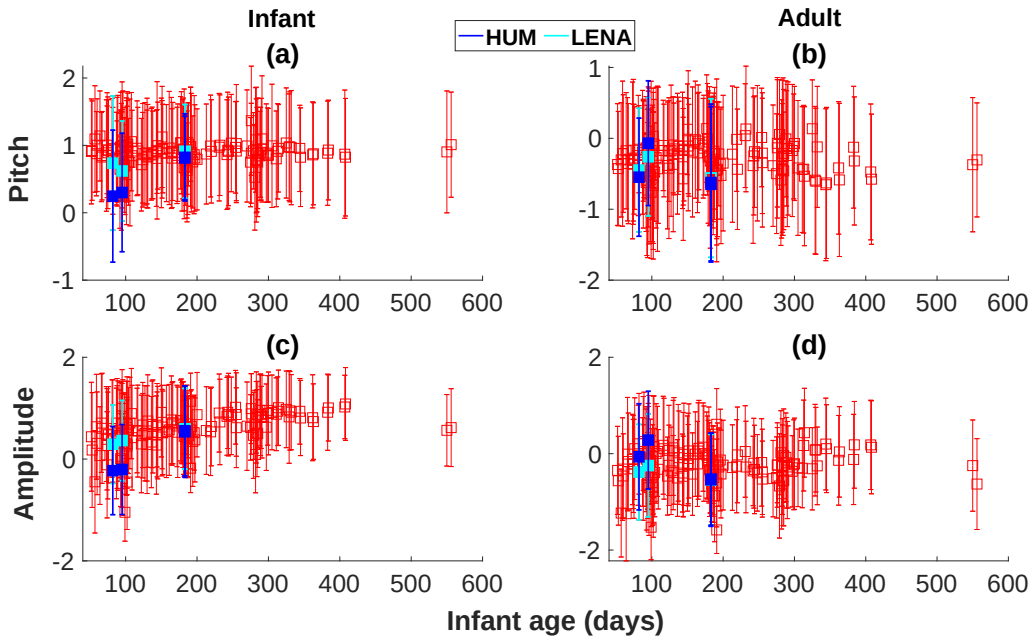

Figure B.1: **Infant and adult vocalisation acoustics as a function of infant age.** Mean pitch (standardised—and hence, unitless—across all infant and adult vocalisations combined) as a function of infant age for infants (panel a) and adults (panel b). Red squares are for the complete dataset, automatically (LENA) labelled. Each data point corresponds to one day-long recording. Error bars represent the standard deviation. Note that we computed the unbiased sample standard deviation, dividing by N-1, where N is the number of samples used in computing the standard deviation. This is in contrast to the population standard deviation, which divides by N. The dark blue points (HUM in the legend) represent the same information as computed from data labelled by human listeners, while the cyan points (LENA in the legend) are the values obtained from the corresponding data as labelled by the LENA system. Note that data from infant 340 at 183 days was re-labelled by two human listeners and results from both listeners are represented. (c, d) Mean standardised amplitude of vocalisations as a function of infant age for infants and adults, respectively.

Table B.2: **Vocalisation acoustics as a function of infant age**. βs are shown with p-values in brackets. Both were obtained from linear mixed effects models with participant ID as a random effect and infant age as fixed effect. Statistically significant results (at a significance level of 0.05) are in bold. All values reported have been rounded to two decimal points wherever possible.

Measure	**Age effect**
Infant mean pitch	0.11 (p=0.13)
Infant pitch standard deviation	**0.21 (p=0.01)**
Infant mean amplitude	**0.50 (p<0.001)**
Infant amplitude standard deviation	**-0.61 (p<0.001)**
Adult mean pitch	-0.12 (p=0.13)
Adult pitch standard deviation	**0.31 (p<0.001)**
Adult mean amplitude	0.001 (p=0.99)
Adult amplitude standard deviation	**-0.35 (p<0.001)**

Table B.3: **Vocalisation acoustics as a function of whether the vocalisation was preceded by response and infant age, with optional response-infant age interaction**. βs are shown with p-values in brackets. Both were obtained from linear mixed effects models with participant ID as a random effect, and whether the preceding vocalisation received a response and infant age as fixed effects. Interaction between response and infant age was used as an optional fixed effect. Results from the model with the interaction term are given in columns 4, 5, and 6. Statistically significant results (at a significance level of 0.05) are in bold. All values reported have been rounded to two decimal points wherever possible.

Measure	**Response**	**Infant age**	**Response**	**Infant age**	**Infant age-response interaction**
Infant mean	**-0.91**	**0.14**	**-0.91**	0.09	0.11
std. dev.	**(p<0.001)**	**(p=0.003)**	**(p<0.001)**	(p=0.19)	(p=0.22)
Infant pitch	-0.14	**0.26**	-0.14	**0.22**	0.08
std. dev.	(p=0.19)	**(p<0.001)**	(p=0.19)	**(p=0.01)**	(p=0.44)
Infant mean	**0.31**	**0.51**	**0.31**	**0.62**	**-0.22**
amplitude	**(p<0.001)**	**(p<0.001)**	**(p<0.001)**	**(p<0.001)**	**(p=0.01)**
Infant amplitude std. dev.	**-0.25**	**-0.54**	**-0.25**	**-0.65**	**0.21**
	(p=0.01)	**(p<0.001)**	**(p=0.01)**	**(p<0.001)**	**(p=0.02)**
Adult mean	**0.80**	**-0.22**	**0.8**	**-0.14**	-0.17
pitch	**(p<0.001)**	**(p<0.001)**	**(p<0.001)**	**(p=0.04)**	(p=0.06)
Adult pitch	**-0.56**	**0.25**	**-0.56**	**0.25**	-0.0002
std. dev.	**(p<0.001)**	**(p<0.001)**	**(p<0.001)**	**(p<0.001)**	(p=1.00)
Adult mean	**1.02**	-0.03	**1.02**	0.04	-0.13
amplitude	**(p<0.001)**	(p=0.49)	**(p<0.001)**	(p=0.50)	(p=0.09)
Adult amplitude std. dev.	**-0.23**	**-0.22**	**-0.23**	**-0.38**	**0.31**
	(p=0.02)	**(p<0.001)**	**(p=0.02)**	**(p<0.001)**	**(p=0.001)**

B.2 Steps in time are correlated with steps in acoustic space

Table B.4: **Step sizes in 2D acoustic space as a function of steps in time**. βs are shown with p-values in brackets. Both were obtained from linear mixed effects models with participant ID as a random effect. Infant age and steps in time were used as fixed effects. Statistically significant results (at a significance level of 0.05) are in bold. All values reported have been rounded to two decimal points wherever possible.

Vocaliser	Time step	Infant age
Adult	**0.05 (p<0.001)**	**0.06 (p<0.001)**
Adult (WR)	**0.11 (p<0.001)**	**0.07 (p<0.001)**
Adult (WOR)	**0.05 (p<0.001)**	**0.06 (p<0.001)**
Infant	**0.07 (p<0.001)**	0.01 (p=0.06)
Infant (WR)	**0.09 (p<0.001)**	**0.02 (p=0.002)**
Infant (WOR)	**0.06 (p<0.001)**	0.001 (p=0.76)

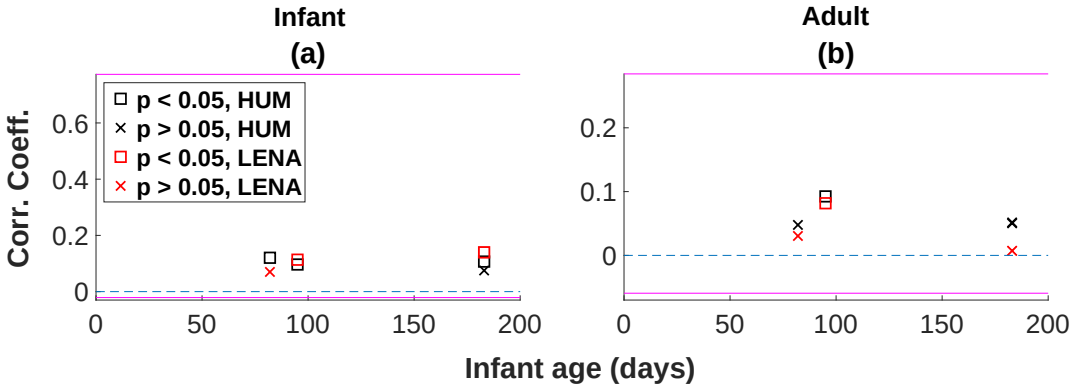

Figure B.2: **Correlations between steps in acoustic space and steps in time validated using data labelled by human listeners.** In (a) and (b), squares represent significant correlations, while crosses represent correlations that were not significant (at a significance level of 0.05), for infant and adult vocalisations, respectively. Black points are obtained from data labelled by human listeners, while red points are from the same data labelled by the LENA software. Note that data from infant 340 at 183 days was re-labelled by two human listeners and results from both listeners are represented. The pink lines indicate the range of space-time correlation values obtained from all available data as labelled by LENA. The dashed blue line separates positive and negative correlation strengths. Note that these correlations were computed using MATLAB's `corrcoef` function.

B.3 Step size probability density fits

For each set of step sizes—infants' pitch steps following an adult response, infants' pitch steps following no adult response, etc.—we used AIC (see `https://github.com/AnneSWarlaumont/infant-vocal-foraging/tree/master/Analyses/AIC_Theo_Rhodes_code` for details) to determine the best fit probability density distribution type and the curve parameters. The types of distributions considered were normal, lognormal, exponential and pareto distributions. For a given type of step, we determined what type of distribution best fit the majority of day-long recordings. We then analysed distribution parameters only from those recordings for which the best fitting distribution also belonged to that of the majority best fit type, for that set of step size distribution fits.

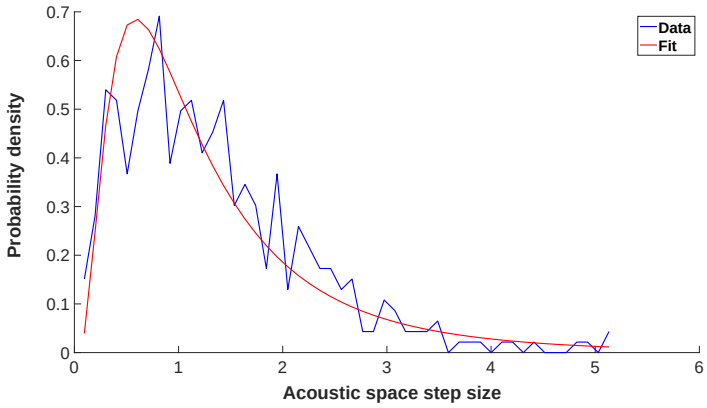

Figure B.3: **Representative example of randomly selected data vs. AIC fit (Infant *mww*, age 75 days).** The figure shows the probability distribution of steps in acoustic space for infant vocalisations where the first infant vocalisation was not followed by an adult response. The distribution derived from the data is in blue and the AIC best fit (lognormal, in this case) is shown in red.

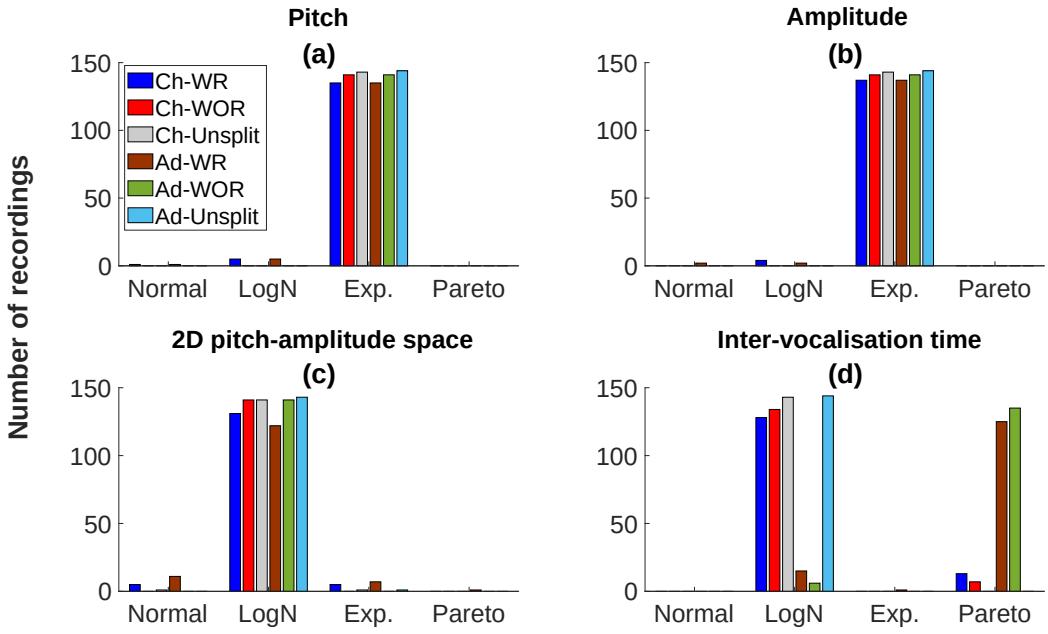

Figure B.4: **Distribution of AIC best fits for step size probability distributions for various step types.** Probability distributions of step sizes along pitch (panel a) and amplitude (panel b) dimensions, respectively, are predominantly exponential, for both adults and infants. (c) Probability distributions of steps in two dimensional acoustic space are predominantly lognormal for both infants and adults. (d) Probability distributions of inter-vocalisation times are predominantly lognormal for infants, and pareto for adults with the exception that unsplit inter-vocalisation time distributions are largely lognormal for adults.

Table B.5: **Goodness of AIC fits.** The means and standard deviations of the R^2 values of the AIC best fit for different step size distribution types are shown. All results are from data labelled by LENA. The step size distributions are organised by whether they were computed from data where the vocaliser was an infant or adult (column 1), and whether they are WOR, WR, or unsplit distributions (column 2). For mean and standard deviation for each step size distribution type for a category (eg. WR pitch step size distributions of adult vocalisations), see `https://osf.io/53amv/`. For a breakdown of the majority best fit for each distribution type, see Fig. B.4. The fifth column contains the mean number of observations per distribution for that category while the sixth column has the total number of distributions that went into calculating the mean and standard deviation of R^2 values for that category. For example, the first row of the table gives the mean and standard deviation of R^2 values of all unsplit step size distributions (pitch, amplitude, 2d acoustic space, and time) where the vocaliser was an infant, regardless of best fit type. For this category, each distribution, on average, had 1043.86 observations, and 572 distributions were used to calculate the mean and standard deviation R^2 values. R^2 values typically fall between 0 to 1, with values closer to 1 indicating better fits. Note that one possible reason for lower R^2 values could be that some step types were less prevalent and therefore had fewer steps on which to fit the distribution (see fifth column of the table). All values reported have been rounded to two decimal points wherever possible. For a similar table for human-labelled data and the corresponding LENA-labelled subset, see Table B.16.

Vocaliser	**Step type**	**Mean R^2**	**Std. dev**	**Mean observations per distribution**	**Number of distributions**
Infant	Unsplit	0.92	0.10	1043.86	572
Infant	WOR	0.91	0.09	603.18	564
Infant	WR	0.78	0.20	184.83	564
Adult	Unsplit	0.94	0.08	3223.40	572
Adult	WOR	0.96	0.04	2031.54	564
Adult	WR	0.60	0.40	137.27	564

B.3.1 Do amplitude step sizes vary with response and infant age?

As shown in Table B.6 and Figure B.5, we observed a significant decrease with age in the 90^{th} percentile value of infant amplitude step size. We also observed a significant increase with age in the λ parameter of the fitted exponential distributions. Both findings suggest that infants take shorter steps in amplitude as they get older, indicating more focused exploration. For adults' amplitude step sizes, we observed a significant increase in the median and 90^{th} percentile values, as well as smaller λ for the exponential fits following infant response and with increasing infant age (Table B.7, Fig. B.5). These findings all point to larger steps in amplitude, and thus suggest broader adult exploration, both with infant response and with increasing infant age.

Finally, Figure B.6 shows how, for infant and adult vocalisations from the entire dataset (all recordings at all ages combined), the median and 90^{th} percentile values of amplitude steps change as a function of the number of vocalisation events by the speaker since a response was last received. Here, the step from the vocalisation that receives a response to the next vocalisation is designated as vocalisation 0, the following step is designated 1, and so on. This continues until the next response is received, at which point the count resets to 0. Thus, as shown in panels (c) and (d), there are fewer steps included in the analysis as the number of events since last response increases, and thus the estimates in panels (a) and (b) can be expected to be less stable as the number of vocalisations since last response increases. At the moment, we treat these visualisations as exploratory, leaving statistical analyses of such multi-event sequences for future work.

Table B.6: **Infant steps in amplitude at the day-long recording level as a function of recent response, infant age, and sample size: results of statistical analyses**. βs are shown with p-values in brackets. Statistically significant results (at a significance level of 0.05) are in bold. Fixed effects are in rows and dependent variables are in columns. A separate linear mixed effects regression model was run for each dependent variable. Sample size was included to control for possible co-variation between sample size (number of vocalisation step events in the recording) and distribution fits, age, and response; including sample size did not affect which results were statistically significant or their sign. Infant ID was a random effect in all models. All values reported have been rounded to two decimal points wherever possible.

		Median	90 th percentile	Exponential parameter, λ
With sample size	Response	-0.13 (p=0.25)	-0.18 (p=0.09)	0.17 (p=0.11)
	Infant age	-0.08 (p=0.17)	**-0.24 (p<0.001)**	**0.15 (p=0.01)**
	Sample size	-0.13 (p=0.12)	-0.12 (p=0.13)	0.10 (p=0.24)
W/o sample size	Response	-0.13 (p=0.25)	-0.18 (p=0.09)	0.17 (p=0.1)
	Infant age	-0.11 (p=0.05)	**-0.27 (p<0.001)**	**0.17 (p=0.002)**

Table B.7: **Adult steps in amplitude at the day-long recording level as a function of recent response, infant age, and sample size: results of statistical analyses**. βs are shown with p-values in brackets. Statistically significant results (at a significance level of 0.05) are in bold. Fixed effects are in rows and dependent variables are in columns. A separate linear mixed effects regression model was run for each dependent variable. Sample size was included to control for possible co-variation between sample size (number of vocalisation step events in the recording) and distribution fits, age, and response; including sample size did not affect which results were statistically significant or their sign. Infant ID was a random effect in all models. All values reported have been rounded to two decimal points wherever possible.

		Median	90 th percentile	Exponential parameter, λ
With sample size	Response	**0.47 (p<0.001)**	**0.37 (p<0.001)**	**-0.43 (p<0.001)**
	Infant age	**0.23 (p<0.001)**	**0.11 (p=0.02)**	**-0.19 (p<0.001)**
	Sample size	0.06 (p=0.37)	0.07 (p=0.24)	0.01 (p=0.92)
W/o sample size	Response	**0.47 (p<0.001)**	**0.37 (p<0.001)**	**-0.43 (p<0.001)**
	Infant age	**0.23 (p<0.001)**	**0.12 (p=0.02)**	**-0.19 (p<0.001)**

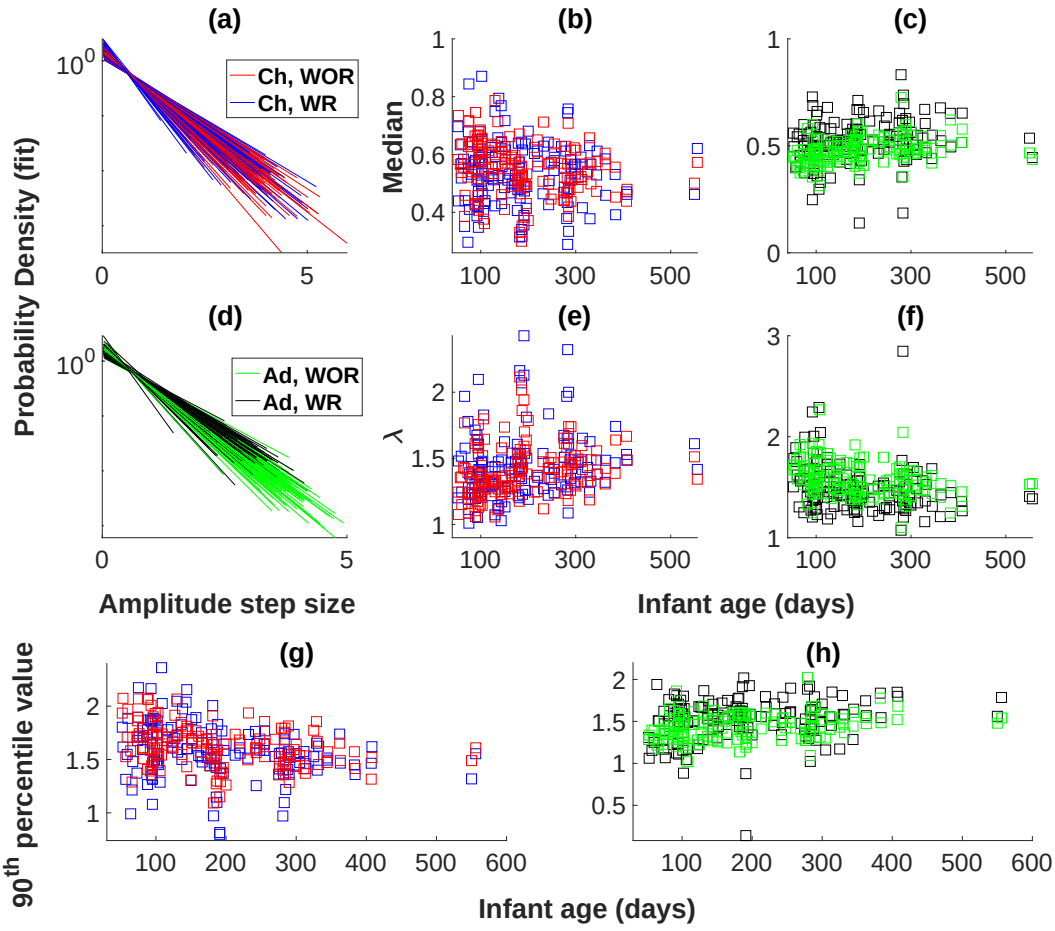

Figure B.5: **Results of amplitude step size distribution analyses.** (a) Exponential probability distributions of infant inter-vocalisation changes ('step sizes') in amplitude following adult response (WR, blue) and not (WOR, red). (b) Median infant amplitude step size as a function of infant age, WR (blue) and WOR (red); (c) shows a similar plot for adults where WR is in black and WOR is in green. (d) Exponential distributions of amplitude steps for adults (WR, black; WOR, green). (e) Age-dependent change in the exponential paramater λ of the WR (blue) and WOR (red) exponential amplitude step size probability distributions for infants; (f) shows a similar plot for adults (WR, black; WOR, green). Note that only distributions that were determined to best fit to an exponential based on AIC are represented in (a), (d), (e), and (f). 90^{th} percentiles of infant amplitude step size distributions plotted against infant age following adult response (WR, blue) and not (WOR, red); (h) shows a similar plot for adult amplitude step sizes (WR, black; WOR, green). Medians and 90^{th} percentile values were computed based on the raw data prior to determining best fits using AIC.

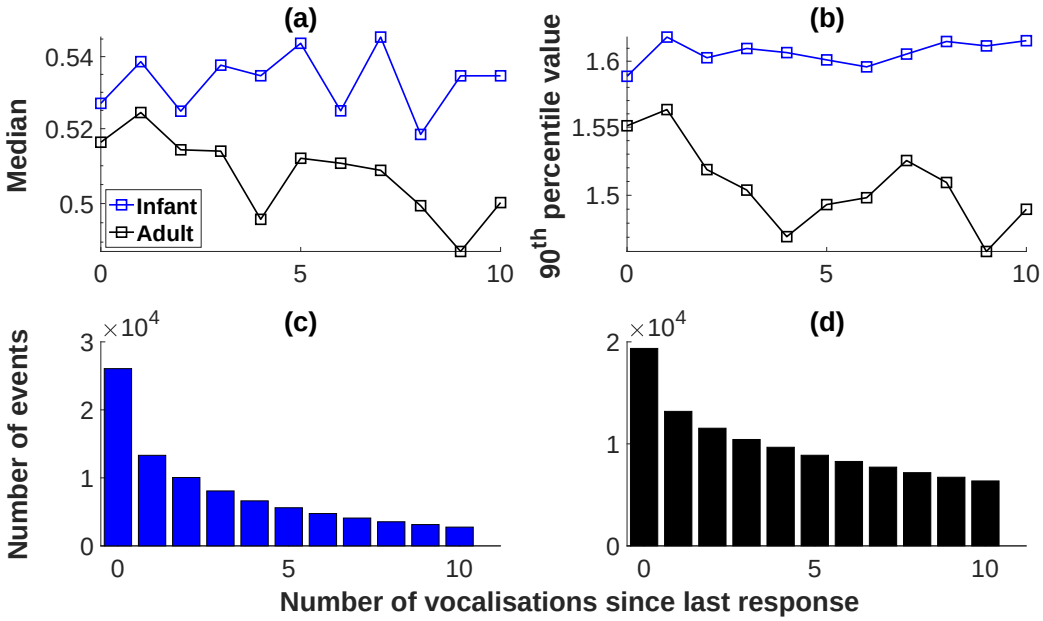

Figure B.6: **Median and 90$^{\text{th}}$ percentile values for amplitude steps as a function of number of events since last response.** (a) Median amplitude step sizes for the entire dataset as a function of number of vocalisations since a response was last received. Here, the step from the vocalisation that received a response to the next vocalisation was assigned 0, the following step (assuming no response to the second vocalisation) was designated 1, and so on. This continued until the next response was received at which point the count resets to 0. The data for infants is in blue and that for adults is in black. (b) 90$^{\text{th}}$ percentile values of amplitude step sizes for the entire dataset, as a function of number of vocalisations since the last response was received (Infant, blue; adult, black). Number of vocalisation events as a function of number of vocalisations since the last response was received for infants (c), and adults (d). Note that plots were terminated after the 10$^{\text{th}}$ vocalisation step following the post-response step.

B.3.2 Do pitch step sizes vary with response and infant age?

Table B.8: **Infant steps in pitch at the day-long recording level as a function of recent response, infant age, and sample size: results of statistical analyses**. βs are shown with p-values in brackets. Statistically significant results (at a significance level of 0.05) are in bold. Fixed effects are in rows and dependent variables are in columns. A separate linear mixed effects regression model was run for each dependent variable. Sample size was included to control for possible co-variation between sample size (number of vocalisation step events in the recording) and distribution fits, age, and response; including sample size did not affect which results were statistically significant or their sign. Infant ID was a random effect in all models. All values reported have been rounded to two decimal points wherever possible.

		Median	90 th percentile	Exponential parameter, λ
With sample size	Response	**0.28 (p=0.01)**	-0.01 (p=0.91)	-0.11 (p=0.27)
	Infant age	**0.25 (p<0.001)**	**0.43 (p<0.001)**	**-0.37 (p<0.001)**
	Sample size	**-0.29 (p<0.001)**	**-0.32 (p<0.001)**	**0.35 (p<0.001)**
W/o sample size	Response	**0.28 (p=0.01)**	-0.01 (p=0.91)	-0.10 (p=0.35)
	Infant age	**0.19 (p<0.001)**	**0.36 (p<0.001)**	**-0.30 (p<0.001)**

Table B.9: **Adult steps in pitch at the day-long recording level as a function of recent response, infant age, and sample size: results of statistical analyses**. βs are shown with p-values in brackets. Statistically significant results (at a significance level of 0.05) are in bold. Fixed effects are in rows and dependent variables are in columns. A separate linear mixed effects regression model was run for each dependent variable. Sample size was included to control for possible co-variation between sample size (number of vocalisation step events in the recording) and distribution fits, age, and response; including sample size did not affect which results were statistically significant or their sign. Infant ID was a random effect in all models. All values reported have been rounded to two decimal points wherever possible.

		Median	90 th percentile	Exponential parameter, λ
With sample size	Response	0.06 (p=0.50)	**-0.19 (p=0.04)**	**0.18 (p=0.03)**
	Infant age	**0.31 (p<0.001)**	**0.32 (p<0.001)**	**-0.40 (p<0.001)**
	Sample size	0.04 (p=0.51)	**0.13 (p=0.04)**	-0.17 (p=0.01)
W/o sample size	Response	0.06 (p=0.50)	**-0.19 (p=0.04)**	**0.17 (p=0.04)**
	Infant age	**0.31 (p<0.001)**	**0.33 (p<0.001)**	**-0.41 (p<0.001)**

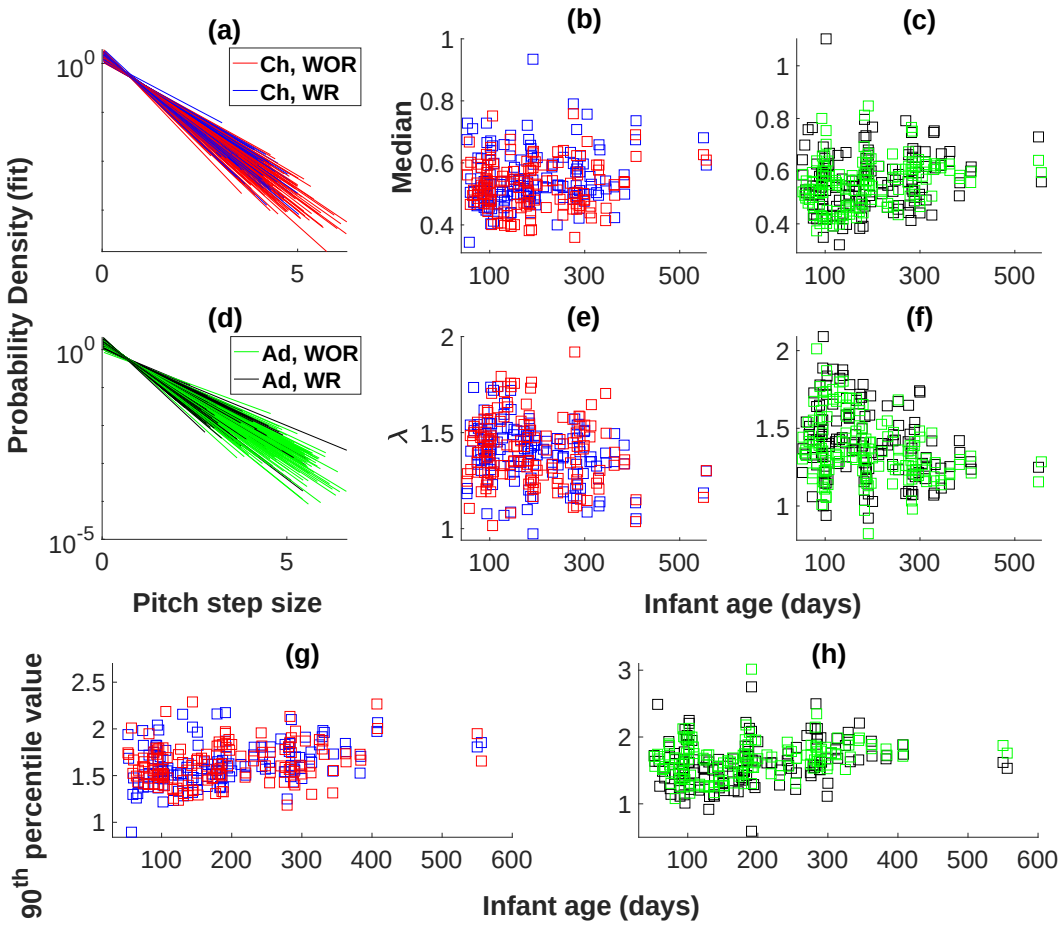

Figure B.7: **Results of pitch step size distribution analyses.** (a) Exponential probability distributions of infant inter-vocalisation changes ('step sizes') in pitch following adult response (WR, blue) and not (WOR, red). (b) Median infant pitch step size as a function of infant age (WR in blue and WOR in red); (c) shows a similar plot for adults where WR is in black and WOR is in green. (d) Exponential distributions of pitch step sizes for adults (WR, black; WOR, green). (e) Age-dependent change in the exponential parameter λ of the WR (blue) and WOR (red) exponential pitch step size probability distributions for infants; (f) shows a similar plot for adults (WR, black; WOR, green). Note that only distributions that were determined to best fit to an exponential based on AIC criterion are represented in (a), (d), (e), and (f). (g) 90^{th} percentiles of the infant pitch step size distributions plotted against infant age following adult response (WR, blue) and not (WOR, red); (h) shows a similar plot for adult pitch step sizes (WR, black; WOR, green). Medians and 90^{th} percentile values were computed based on the raw data, prior to determining best fits using AIC.

For infants, as shown in Table B.8 and Figure B.7, we found a significant increase following adult response in the median pitch step size, suggesting that infants are more likely to take longer steps in the pitch dimension immediately after receiving adult responses. As infant age increased, median and 90^{th} percentile pitch step sizes increased and λs of the exponential step size probability density fits decreased; these three findings suggest that as infants get older, they explore more broadly in the pitch dimension.

For adults, as shown in Table B.9 and Figure B.7, we found a significant decrease in 90^{th} percentile value and a significant increase in λ parameter of the exponential fits following an infant response. Together, these results suggest that adults are more likely to take shorter steps following a response from an infant, indicating more focused exploration. We also found that median and 90^{th} percentile values increased and λ decreased with infant age. These findings suggest that adults take longer steps in the pitch dimension as infant age increases, suggesting more adult pitch exploration and variation as the infant develops.

Finally, Figure B.8 shows the median and 90^{th} precentile values of pitch steps as a function of the number of vocalisation events by the vocaliser since a response was last received, for infant and adult vocalisations from the entire dataset (all recordings at all ages combined).

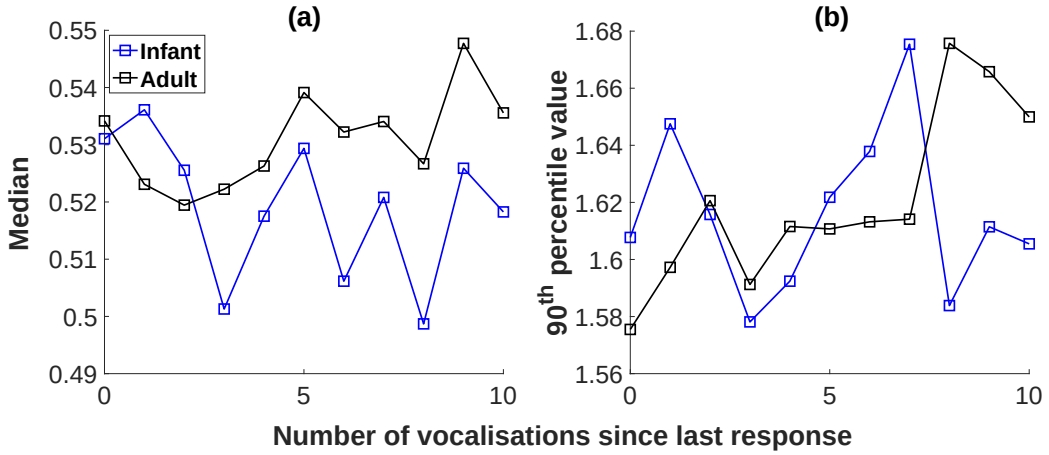

Figure B.8: **Median and 90$^{\text{th}}$ percentile values for pitch steps as a function of number of steps since a response was last received.** (a) Median pitch step size for the entire dataset, as a function of number of vocalisations since a response was last received. The step from the vocalisation that received a response was assigned 0, the following step (assuming no response in the meantime) was assigned 1, and so on. This continued until the next response was received, at which point the count reset to 0. Infant data are shown in blue, and adult data are in black. (b) 90$^{\text{th}}$ percentile of pitch step sizes for the entire dataset, as a function of number of vocalisations since the last response was received (Infant, blue; adult, black). Note that plots were terminated after the 10$^{\text{th}}$ vocalisation event after the last response.

B.3.3 Do step sizes in 2D acoustic space vary with response and infant age?

See Sections 3.2–3.3 for additional results and discussion.

For a demonstration of how lognormal and pareto distributions change as a function of their parameters, see `https://osf.io/2fuje/` (Wolfram Player may be used to view the demo). To see how the variation of the parameters of lognormal and pareto distributions in parameters regimes seen in our data per AIC best fits, see Fig. B.9.

Table B.10: **Steps in 2D acoustic space at the day-long recording level as a function of recent response and infant age: additional results of statistical analysis.** βs are shown with p-values in parentheses. Statistically significant results (at a significance level of 0.05) are in bold. Fixed effects are in rows and dependent variables are in columns. A separate linear mixed effects regression was run for each dependent variable. Infant ID was a random effect in all models. In Tables 3.1–3.2, sample size was included to control for possible co-variation between sample size and distribution fits, age, and response. Results when not controlling for sample size are shown here. The only difference in statistically significant relationships was observed for median infant step size increasing with infant age; this relationship only reached statistical significance when sample size was controlled. All values reported have been rounded to two decimal points wherever possible.

		Median	90 th percentile	Lognormal parameter, μ	Lognormal parameter, σ
		Infant vocalisations			
W/o sample size	Response	-0.02 (p=0.85)	-0.18 (p=0.14)	-0.01 (p=0.95)	**-0.40 (p<0.001)**
	Infant age	0.08 (p=0.16)	0.02 (p=0.75)	0.05 (p=0.40)	0.01 (p=0.80)
		Adult vocalisations			
W/o sample size	Response	**0.24 (p=0.01)**	0.02 (p=0.83)	**0.26 (p=0.003)**	**-0.60 (p<0.001)**
	Infant age	**0.39 (p<0.001)**	**0.30 (p<0.001)**	**0.46 (p<0.001)**	-0.15 (p=0.01)

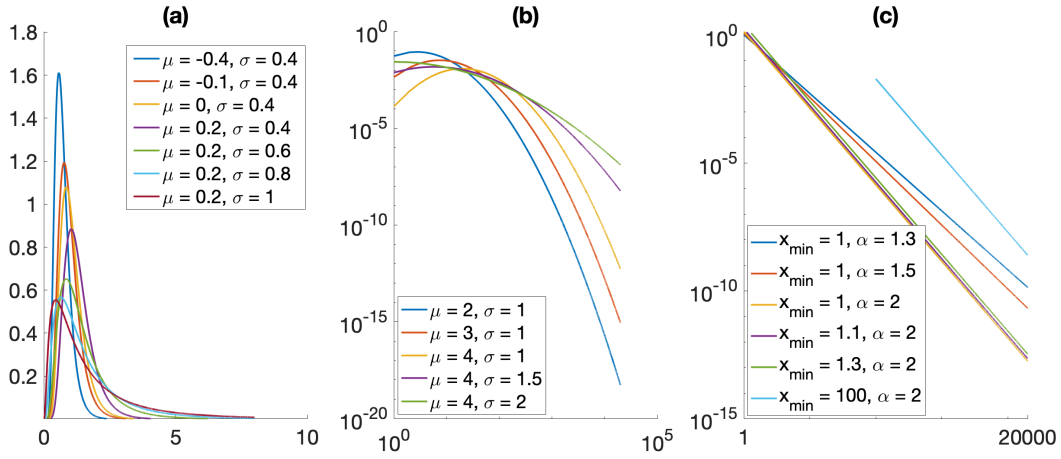

Figure B.9: **Representative probability distributions to demonstrate how values affect fitted step size distribution shapes.** (a) has lognormal probability distributions for parameter ranges similar to those observed in AIC fits for infant and adult WR and WOR step size distributions in 2D acoustic space obtained from LENA and human-labelled data. As μ increases the peak widens while shifting to the right and the tail gets wider, making both intermediate steps and larger steps more likely. As σ increases, the peak widens while shifting to the left, and the tail widens, making shorter and larger steps more likely. (b) A similar plot for lognormal parameter ranges similar to those observed in AIC fits for WR and WOR temporal step size distributions (infants) obtained from LENA and human-labelled data. In contrast with plot (a), plot (b) uses log scales for both x and y axes, in order to better highlight the effects of differing parameter values in the ranges of interest. (c) A similar plot for pareto parameter ranges similar to those observed in AIC fits for adult WR and WOR temporal step size distributions obtained from LENA and human-labelled data. In addition, we add a reference curve at $x_{min} = 100$ to show the effect of increasing x_{min}. For the range of x_{min} values obtained from AIC best fits ($\sim$ 1-1.3) and the range of step sizes in time present in our data, the change in x_{min} has no appreciable effect on the pareto distribution. In contrast, however, as α increases, the distribution decays rapidly and the likelihood of larger step sizes decrease. For parameter ranges for infant and adult WR/WOR lognormal fits of step size distributions in 2D acoustic space from LENA-labelled data, see Fig. B.10; for parameter ranges for infant and adult WR/WOR step size distributions in time from LENA-labelled data, see Fig. B.12; for parameter values for human-labelled data and the corresponding LENA-labelled subset, see `https://osf.io/xptv4/` and `https://osf.io/56gx9/`. Note that the range of X-axis values used in (a), (b), and (c) correspond to the range of step size values for the data that went into the analyses reflected in Fig. B.10 and B.12. All probability distributions shown have been normalised such that the area under the curve from 0 to the maximum X-axis value shown, is 1.

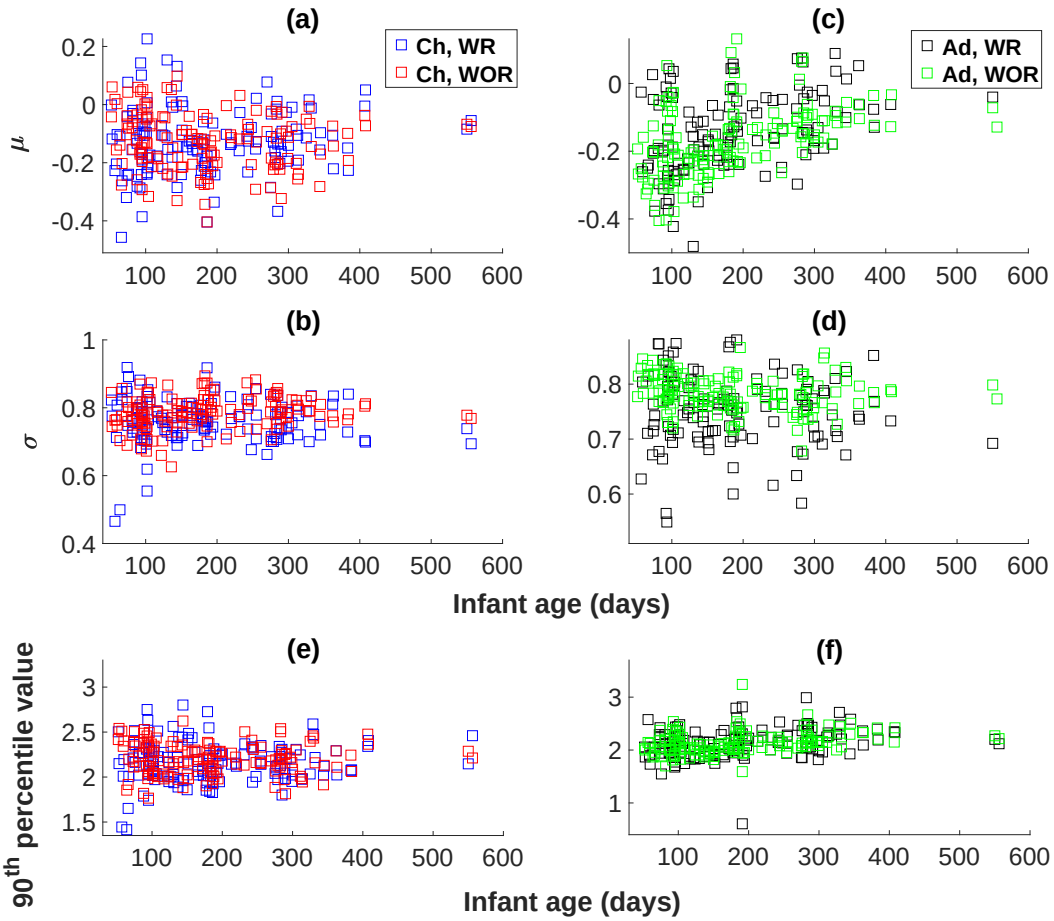

Figure B.10: **Additional results of 2D acoustic space step size distribution analyses.** (a) Infant μ and (b) σ as a function of age and whether the child's first vocalisation received an adult response (WR, blue) or not (WOR, red). (c and d) show similar plots for adult (WR, black; WOR, green). Only distributions that were determined to best fit to lognormal based on AIC criterion are represented in (a-d). (e) 90^{th} percentile values for step sizes in 2D acoustic space as a function of infant age for infants (WR, blue; WOR, red). (f) shows a similar plot for adults (WR, black; WOR, green). 90^{th} percentile values were computed from the raw data, before finding the AIC best fit.

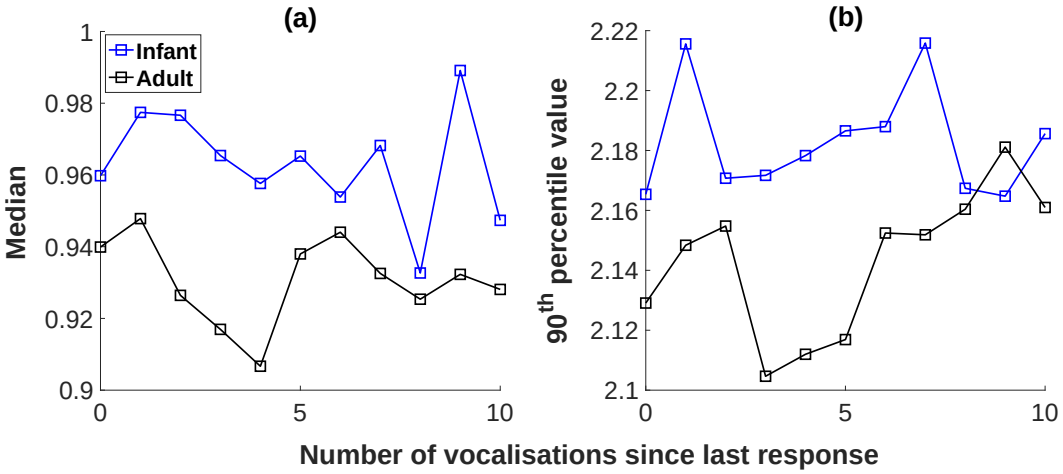

Figure B.11: **Median and 90$^{\text{th}}$ percentile values for steps in 2D acoustic space as a function of number of events since last response.** (a) Median step size in 2D acoustic space for the entire dataset as a function of number of vocalisations since the last response was received. The step from the vocalisation that received a response was assigned 0, the following step (assuming no response in the meantime) was assigned 1, and so on. This continued until the next response was received, at which point the count reset to 0. The data for infants are in blue and data for adults are in black. (b) 90$^{\text{th}}$ percentile value of step sizes in 2D acoustic space for the entire dataset, as a function of number of vocalisations since the last response was received (infant, blue; adult, black). Note that plots were terminated after the 10$^{\text{th}}$ vocalisation event after the last response.

B.3.4 Do inter-vocalisation intervals vary with response and infant age?

See Sections 3.2–3.3 for additional figures and results.

Table B.11: **Inter-vocalisation interval distributions at the day-long recording level as a function of recent response and infant age: results of statistical analysis.** βs are shown with p-values in parentheses. Statistically significant results (at a significance level of 0.05) are in bold. Fixed effects are in rows and dependent variables are in columns. A separate linear mixed effects regression was run for each dependent variable. Infant ID was a random effect in all models. In Tables 3.1–3.2, sample size was included to control for possible co-variation between sample size and distribution fits, age, and response. Here, results when not controlling for sample size are shown. When sample size was excluded, infant age showed a statistically significant negative relationship with the 90th percentile value and lognormal parameters σ and μ of infant inter-vocalisation interval distributions. Similarly, when sample size was included, the positive relationship between adult median inter-vocalisation interval and an infant response being recently received was only marginally significant. Finally, when sample size was included, we found a statistically significant negative effect in the μ parameter of infant inter-vocalisation step size distributions with respect to having recently received an adult response. All values reported have been rounded to two decimal points wherever possible.

		Median	90 th percentile	Lognormal parameter, μ	Lognormal parameter, σ
		Infant vocalisations			
W/o sample size	Response	**-0.25 (p=0.02)**	-0.11 (p=0.29)	-0.18 (p=0.07)	**-0.42 (p<0.001)**
	Infant age	-0.07 (p=0.17)	**-0.13 (p=0.02)**	**-0.15 (p=0.004)**	**-0.17 (p=0.001)**
		Median	90 th percentile	Pareto parameter, x_{min}	Pareto parameter, α
		Adult vocalisations			
W/o sample size	Response	0.16 (p=0.07)	-0.07 (p=0.50)	**0.64 (p<0.001)**	**0.33 (p<0.001)**
	Infant age	**0.19 (p<0.001)**	0.01 (p=0.87)	-0.003 (p=0.96)	**-0.21 (p<0.001)**

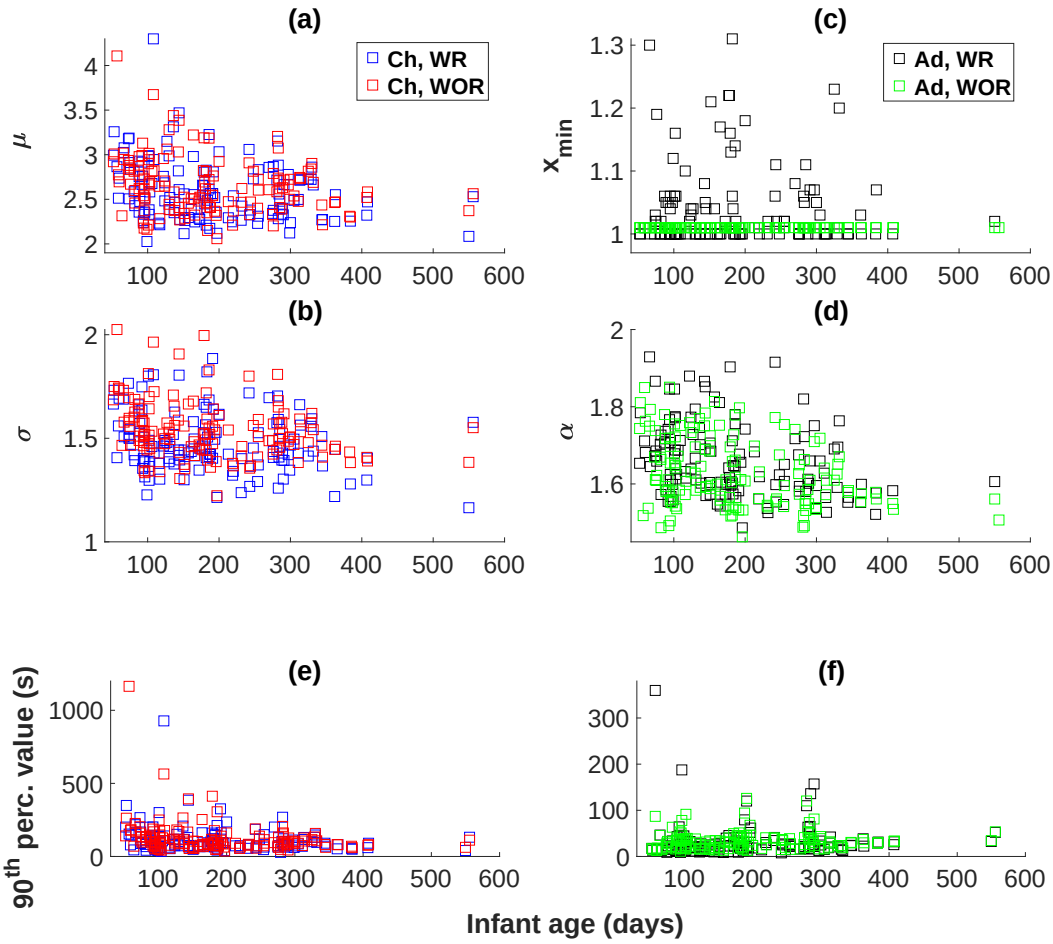

Figure B.12: **Additional results of inter-vocalisation step size distribution analyses.** Infant μ (a) and σ (b) plotted against infant age following adult response (WR, blue) and not (WOR, red). Only distributions that were best fit to lognormal curves based on AIC are represented. Adult x_{min} (c) and α (d) plotted against infant age following infant response (WR, black) and not (WOR, green). Only distributions that were best fit to pareto curves per AIC are shown. (e) shows 90^{th} percentile values for inter-vocalisation intervals plotted against infant age for infants (WR, blue; WOR, red). (f) shows a similar plot for adults (WR, black; WOR, green). 90^{th} percentile values were computed from the raw data, before finding the AIC best fit.

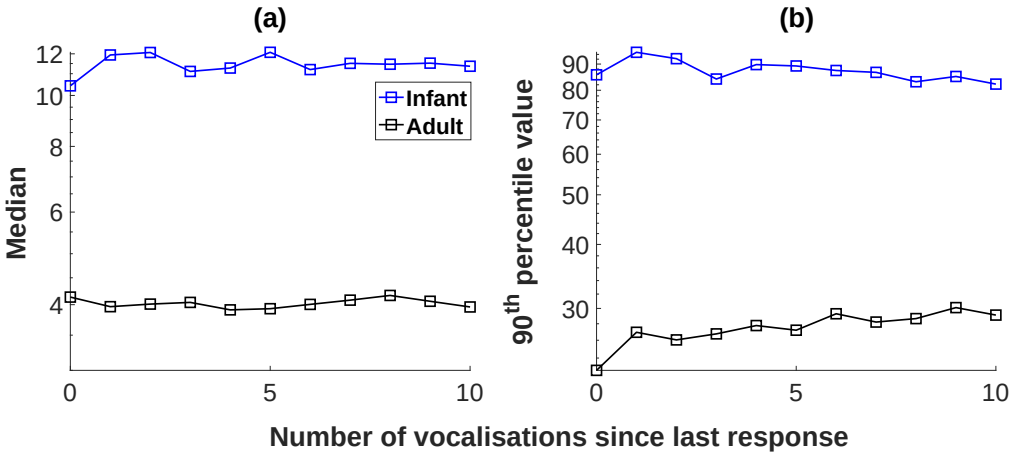

Figure B.13: **Median and 90th percentile values for inter-vocalisation intervals as a function of number of events since a response was last received.** (a) Median inter-vocalisation interval for the entire dataset as a function of number of vocalisations since the last response was received. The step from the vocalisation that received a response to the next vocalisation was assigned 0, the following step (assuming no response in the meantime) was assigned 1, and so on. This continued until the next response was received at which point the count reset to 0. Data for infants are in blue and adult data are in black. (b) 90th percentile value of inter-vocalisation intervals for the entire dataset, as a function of number of vocalisations since the last response was received. Infant data is in blue and adult data is in black. Note that plots are terminated after the 10th vocalisation event after the last response.

Note that unlike other measures presented in this study (mean and standard deviation of pitch and standard deviation, median and 90th percentile values of step sizes in amplitude, pitch, and 2D acoustic space), infant inter-vocalisation intervals are longer than those of adults, based on median and 90th percentile values, i.e., infants vocalise more sparingly than adults. This observation is supported by Table B.5. This disparity in median and 90th percentile values is in contrast to all other measures presented, which are comparable for both infants and adults.

B.4 Using data re-labelled by human listeners to check the validity of automatically labelled data

B.4.1 Inter-rater reliability measures

[Confusion matrices for LENA vs. human-listener labelled data]**Confusion matrices for LENA vs. human-listener labelled data.** Confusion matrices for each set of human-labelled data are shown, with LENA labels as the known class (row indices) and human-listener labels as the predicted class (column indices). For a description of human-listener vs. LENA labels, refer to Section 3.4.

	CHN (H)	FAN (H)	MAN (H)
L1, 274			
CHNSP (L)	320	32	0
FAN (L)	99	295	12
MAN (L)	22	26	83
L2, 530			
CHNSP (L)	954	195	12
FAN (L)	278	445	113
MAN (L)	66	12	287
L1, 340			
CHNSP (L)	462	12	17
FAN (L)	23	269	137
MAN (L)	3	17	512
L3, 340			
CHNSP (L)	502	12	16
FAN (L)	24	258	147
MAN (L)	3	15	541

Table B.12: **Confusion matrix for human listeners 1 vs. 3 for data from infant 340.** Confusion matrix for data from infant 340 labelled by listeners 1 and 3 is shown, with L1 labels as the known class (row indices) and L3 labels as the predicted class (column indices). For a description of human-listener labels, refer to Section 3.4. Note that we see high agreement between L1 and L3 labels, which is in agreement with the high inter-rater reliability scores for listeners 1 and 3 for data from infant 340 (see Table 3.4).

Infant 340	CHN (L3)	FAN (L3)	MAN (L3)
CHN (L1)	478	0	0
FAN (L1)	0	266	10
MAN (L1)	1	0	631

Table B.13: **Fraction of LENA speaker labels as identified by human listeners.** The fraction of infant speech-related (CHNSP) labels and adult (MAN or FAN) labels as labelled by LENA that fall into various label categories as determined by human listeners is reported, for each human-labelled dataset. The human-labelled categories include infant speaker label (CHN) only (Ch only, column 2); adult (MAN or FAN or both) speaker only (Ad only, column 3); infant speaker with possible other child overlap but no adult overlap (column 4); adult speaker with possible other child (not the infant wearing the recorder) overlap (column 5); infant (CHN) and/or adult (MAN or FAN) speakers with possible overlap of other child voice (column 6); and either no voices identified or other child identified, but no target infant (CHN) or adult (MAN or FAN). For a description of human-listener vs. LENA labels, refer to the Section 3.4. All values reported have been rounded to two decimal points wherever possible. We see fairly good agreement ($\sim$ 70% or greater) between human listeners and LENA for the infant 340 data labelled by listeners 1 and 3 both when overlaps are allowed (Adult (L) vs. column 5, and Ch (L) vs. column 4) and not allowed (Adult (L) vs. column 3, and Ch (L) vs. column 2). Note that these are also the human-labelled datasets with the highest percent agreement and Cohen's Kappa values (Table 3.4). These values are lower for data from infant 274 labelled by listener 1, and data from infant 530 labelled by listener 2, both of which have lower inter-rater reliability scores. In addition, the data loss as a consequence of more restrictions on determining what is an infant vocalisation (CHN only vs. CHN+Olp; no Ad vs. Ch/Ad) or an adult vocalisation (Ad only vs. Ad+Olp; no Ch vs. Ch/Ad) for human-labelled data can be seen clearly in the table.

	Ch only (H)	Ad only (H)	Ch+Olp; no Ad (H)	Ad+Olp; no Ch (H)	Ch/Ad (H)	No Ch/Ad (H)
L1, 274						
Ch (L)	0.49	0.05	0.61	0.10	0.83	0.17
Ad (L)	0.08	0.32	0.11	0.45	0.87	0.13
L1, 340						
Ch (L)	0.67	0.04	0.72	0.08	0.86	0.14
Ad (L)	0.02	0.74	0.02	0.89	0.96	0.04
L3, 340						
Ch (L)	0.72	0.04	0.78	0.06	0.90	0.09
Ad (L)	0.02	0.75	0.02	0.87	0.96	0.04
L2, 530						
Ch (L)	0.51	0.11	0.57	0.15	0.83	0.17
Ad (L)	0.14	0.37	0.16	0.43	0.74	0.26

B.4.2 Acoustic space trajectories and step size distributions of infants and adults

In Figures B.14, B.15, B.16, and B.17, we present data from three infants at different ages re-labelled by human listeners. For comparison, we also present the corresponding data (i.e. from the same day-long recording) as labelled by the LENA software. In addition to comparing how data labelled by the LENA software compares to the same data labelled by human listeners, we have one recording (participant 340, age 183 days) labelled by two different human listeners to compare how differences in labelling by different listeners affect the data.

In each of the four figures, panel (a) shows infant vocalisations' locations in 2-D acoustic space and panel (b) shows adult vocalisations' locations in the same space. The data are depicted as a series of directed vectors from vocalisation i at a location in the acoustic space given by the ordered pair (f_i, d_i) to vocalisation $i+1$ at (f_{i+1}, d_{i+1}), starting from the first available vocalisation based on the data. Here, f is the z-scored log pitch, and d is the z-scored amplitude. Pink vectors start at infant vocalisations that have received adult responses (corresponding to WR steps) and grey vectors start at infant vocalisations that have not received responses (corresponding to WOR steps), as labelled by the human listener. These plots are on the left side of the (a) panels. On the left side of panel (b), pink vectors start at adult vocalisations that have received infant responses (WR steps) and grey vectors start at adult vocalisations that have not received infant responses (WOR steps), as labelled by the human listener. On the right side of panel (a) are infant step data based on LENA labels, with blue vectors starting at infant vocalisations that have received adult responses (WR steps) and red vectors starting at infant vocalisations that have not received responses (corresponding to WOR steps). Finally, on the right side of panel (b) are the adult step data based on LENA labels, with black vectors starting at adult vocalisations that have received infant responses (WR steps) and green vectors starting at adult vocalisations that have not received responses (WOR steps). Note that vocalisations that were marked as 'not applicable' for whether a response was received or not were excluded from these plots.

Finally, we present raw and fitted probability distributions of steps in pitch, amplitude, 2D acoustic space, and time. Each plot shows a set of human-labelled data together with the same recording's data as labelled by the LENA software. Plots c1–c4 and d1–d4 show raw probability distributions of step sizes following response (WR) based on human labelled utterances (pink) and the corresponding data as labelled by LENA (infant: blue; adult:

black). Plots c5–c8 and d5–d8 show raw probability distributions of step sizes following no response (WOR) for human labelled utterances (grey) and the corresponding data labelled by LENA (infant: red; adult: green). Plots e1–f8 show best fit curves for step size probability distributions from human labelled data and the corresponding data labelled by LENA. Dashed pink lines are for human-labelled data WR (where the first vocalisation received a response), dashed grey are for human-labelled data WOR (where the first vocalisation did not receive a response), dashed blue lines are for the corresponding LENA-labelled WR data for infants, and dashed black lines are for the corresponding LENA-labelled adult WR data. Similarly, dashed red lines are LENA-labelled infant WOR data and dashed green lines are LENA-labelled adult WOR data. More details about the fits and fit parameters available here: `https://osf.io/8ern6/`, in folder Reported_and_auxiliary_results.

We see that the without response (WOR) steps' distributions are extremely similar when LENA's labels vs. human re-labelling are used. On the other hand, for with-response steps, we see larger differences between the distribution fits (shown in dashed lines in plots e1–e4 and f1–f4) for human-labelled vs. LENA-labelled data. We also see that the raw distributions are much less smooth for human-labelled with-response data. A likely reason for these discrepancies is the paucity of with-response data in the human-labelled dataset (see Table B.14). This in turn may perhaps be due to human listeners' greater sensitivity to voices, so that segments labelled by humans may have been more likely to be identified as containing multiple human voices and therefore excluded from analysis. Additionally, exclusion of many of segments of the audio (based on LENA's automatic segmentation and classification) from the human labelling task made it unlikely that relevant voices that were missed by the LENA system would have been included in the human analysis. For all these reasons, responses were less likely to be identified within the human-labelled datasets. For quantitative results on how step size distributions based on human-labelled data compare to those based in corresponing LENA-labelled data, see Table B.15.

Table B.14: **Number of data points in WR and WOR step sizes for human re-labelled data and corresponding LENA data.** The number of data points for both WR and WOR step sizes for all three datasets that were re-labelled by human listeners are shown. These numbers are reported for adult vocalisations (Ad) and infant vocalisations (Ch), for both human re-labelled data (HUM) and the corresponding LENA data (LENA). The LENA-labelled data include many more steps for all types except child steps in which the first vocalisation was not followed by an adult response. We find that the number of steps for this category is comparable for both human-labelled data and corresponding LENA-labelled data.

Child ID, age	LENA Ad		LENA Ch		Listener ID	HUM Ad		HUM Ch	
	WR	WOR	WR	WOR		WR	WOR	WR	WOR
274, 82 days	100	941	124	383	L1	7	359	10	338
340, 183 days	76	1001	86	438	L1	7	771	10	397
340, 183 days	76	1001	86	438	L3	15	781	13	425
530, 95 days	295	1769	340	1289	L2	7	492	17	1060

Table B.15: **Two-sample Kolmogorov-Smirnov (KS2) test results, comparing step size distributions from human-labelled data and corresponding LENA labelled data.** f_{rej} represents the fraction of tests which failed to reject the null hypothesis (that data from the two samples – step size distriution from data labelled by LENA and human listeners – are drawn from the same distribution) at the 0.05 significance level, for individual datasets (column 1). For each category specified in columns 3, 5, and 7, we present the mean p-value from the KS2 test in columns 4, 6, and 8, respectively. The standard deviation of the p-value in parentheses. For example, for data from infant 274 at 82 days labelled by listener 1, 83 percent of KS2 tests (performed on unsplit, WR, and WOR step size distributions in pitch, amplitude, 2D acoustic space, and time, for infants and adults) failed to reject the null hypothesis. Further, for all unsplit step size distributions (pitch, amplitude, 2D acoustic space, and time) where the infant was the vocaliser, the mean p-value associated with the KS2 tests performed was 0.37, with a standard deviation of 0.35. By and large, we see that f_{rej} is high for all datasets except infant 530 at 95 days labelled by listener 2. Note that f_{rej} is lowest for data from infant 530 at 95 days labelled by listener 2, followed by data from infant 274 at 82 days labelled by listener 1, both of which have the lowest reliability scores (Table 3.4; Subsection B.4.1). Similarly, we also see the lowest mean p-values for these datasets. All values reported have been rounded to two decimal points wherever possible. For more detailed results, see `https://osf.io/8ern6/` .

Infant ID, age, and listener ID	f_{rej}						
Infant 274, 82 days, listener L1	0.83	Infant (unsplit)	0.37 (0.35)	Infant (WR)	0.33 (0.32)	Infant (WOR)	0.46 (0.29)
		Adult (unsplit)	0.18 (0.13)	Adult (WR)	0.51 (0.39)	Adult (WOR)	0.28 (0.19)
Infant 340, 183 days, listener L1	0.92	Infant (unsplit)	0.35 (0.35)	Infant (WR)	0.79 (0.13)	Infant (WOR)	0.39 (0.39)
		Adult (unsplit)	0.74 (0.49)	Adult (WR)	0.76 (0.36)	Adult (WOR)	0.73 (0.49)
Infant 340, 183 days, listener L3	0.92	Infant (unsplit)	0.69 (0.22)	Infant (WR)	0.65 (0.12)	Infant (WOR)	0.70 (0.31)
		Adult (unsplit)	0.67 (0.45)	Adult (WR)	0.63 (0.33)	Adult (WOR)	0.69 (0.46)
Infant 530, 95 days, listener L2	0.46	Infant (unsplit)	<0.001 (<0.001)	Infant (WR)	0.51 (0.36)	Infant (WOR)	0.02 (0.02)
		Adult (unsplit)	0.05 (0.05)	Adult (WR)	0.65 (0.35)	Adult (WOR)	0.25 (0.42)

Table B.16: **Goodness of AIC fits (human-labelled data and corresponding LENA-labelled subset).** The means and standard deviations of the R^2 value of the AIC best fit for different step size distribution types are shown. All results are from data labelled by human listeners and the corresponding LENA-labelled subset. The step size distributions are organised by whether they were computed from data where the vocaliser was an infant or adult (column 1), and whether they are WOR, WR, or unsplit distributions (column 2). For mean and standard deviation for R^2 values for each step size distribution type for a category (eg. WR pitch step size distributions of adult vocalisations), see `https://osf.io/53amv/`. For a breakdown of the majority best fit for each distribution type, see Fig. B.4. The sixth column has the mean number of observations per distribution for that category while the seventh column has the total number of distributions that went into calculating the mean and standard deviation of R^2 values for that category. For example, the first row of the table gives the mean and standard deviation of all unsplit step size distributions (pitch, amplitude, 2D acoustic space, and time) where the vocaliser was an infant as labelled by human listeners, regardless of best fit type. For this category, each distribution on average had 684 observations, and 16 distributions were used to calculate the mean and standard deviation R^2 values. All values reported have been rounded to two decimal points wherever possible. R^2 values typically fall between 0 to 1, with values closer to 1 indicating better fits. Note that one possible reason for lower R^2 values could be that some step types were less prevalent and therefore had fewer steps on which to fit the distribution (see the sixth column of the table).

Labelled by	**Vocaliser**	**Step type**	**Mean R^2**	**Std. dev**	**Mean num. of observations**	**Num. of distributions**
HUM	Infant	Unsplit	0.90	0.06	684	16
LENA (subset)	Infant	Unsplit	0.94	0.03	1188.67	12
HUM	Infant	WOR	0.90	0.09	555	16
LENA (subset)	Infant	WOR	0.94	0.04	703.33	12
HUM	Infant	WR	0.22	0.24	12.5	16
LENA (subset)	Infant	WR	0.69	0.21	183.33	12
HUM	Adult	Unsplit	0.90	0.06	861.5	16
LENA (subset)	Adult	Unsplit	0.93	0.03	1930.33	12
HUM	Adult	WOR	0.90	0.09	600.75	16
LENA (subset)	Adult	WOR	0.94	0.06	1237	12
HUM	Adult	WR	0.18	0.24	9	16
LENA (subset)	Adult	WR	0.69	0.22	157	12

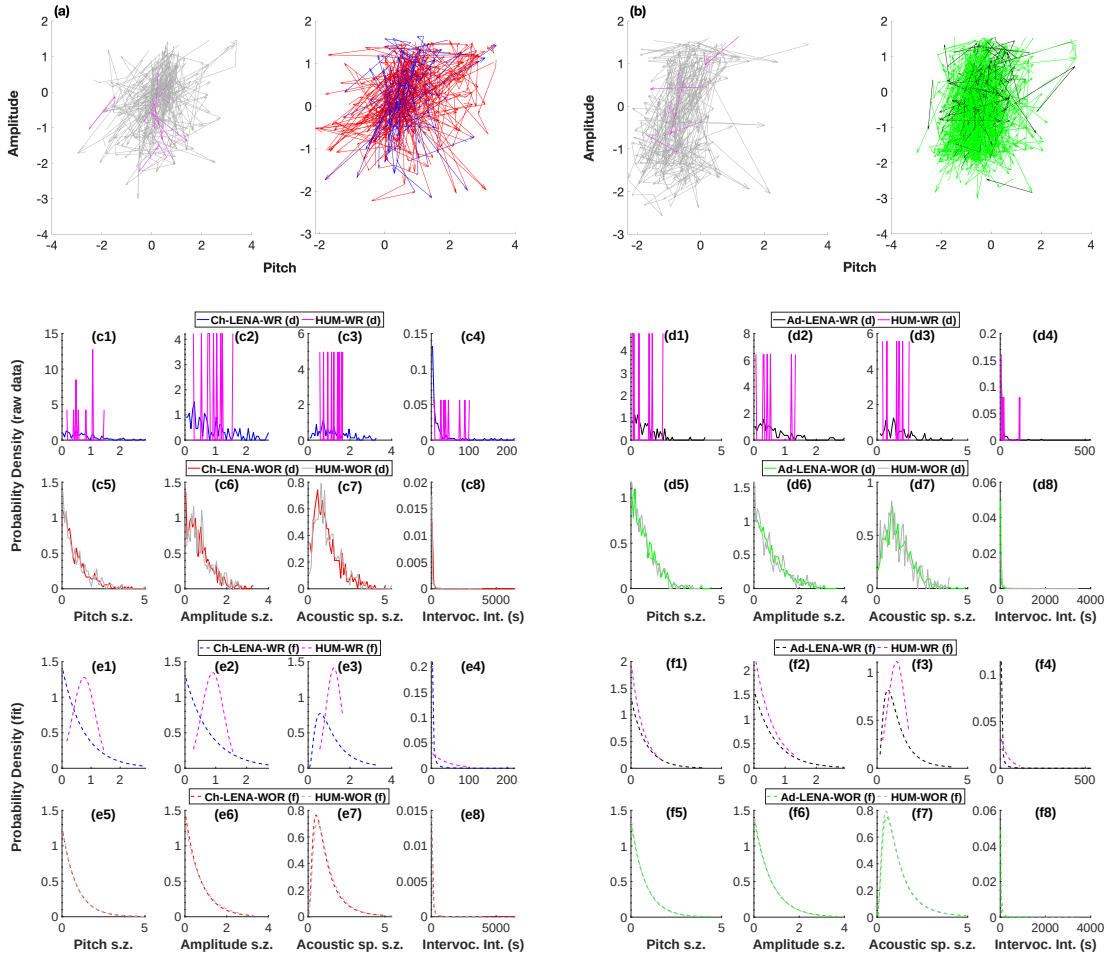

Figure B.14: **Human labelled data and corresponding LENA labelled data - Participant 274 at 82 days old; Listener 1.** (a) Acoustic space traversed by the infant as labelled by human listener L1 (WR in pink and WOR in grey; left), and as labelled by the LENA software (WR in blue and WOR in red; right). (b) Acoustic space traversed by the adult as labelled by human listener L1 (WR in pink and WOR in grey; left), and as labelled by the LENA software (WR in black and WOR in green; right). Raw and fitted probability distributions of step sizes from human labelled data and corresponding LENA labelled data are shown in panels (c1) through (f8). Infant data for step size distributions following a response (WR) are in panels c1–c4 and infant data for WOR are in panels c5–c8. Adult WR data are in panels d1–d4, and adult WOR data are in panels d5–d8. Fits for Infant WR data are in panels e1–e4 and infant WOR fits are in panels e5–e8. Adult WR fits are in panels f1–f4, and adult WOR fits are in panels f5–f8. The (f) in the legend indicates fitted as opposed to raw data (indicated by (d) in the legend). Note that both human-labelled infant and adult WR and WOR data/fits are given by pink and grey solid/dashed lines and are presented in the same subplots as their corresponding LENA-labelled data/fits, to allow for visual comparison of how well the curves do, or in a few cases do not, overlap.

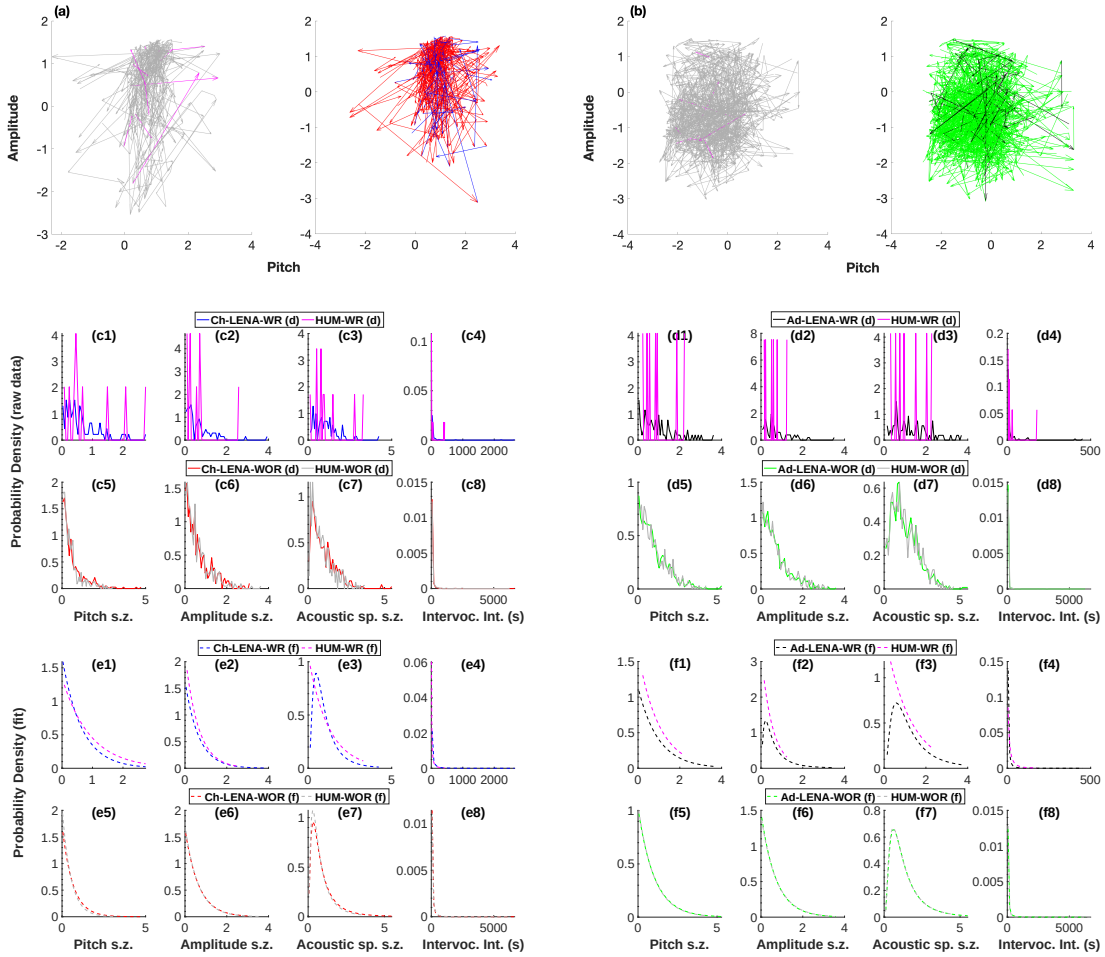

Figure B.15: **Human labelled data and corresponding LENA labelled data - Participant 340 at 183 days old; Listener 1.** (a) Acoustic space traversed by the infant as labelled by human listener L1 (WR in pink and WOR in grey; left), and as labelled by the LENA software (WR in blue and WOR in red; right). (b) Acoustic space traversed by the adult as labelled by human listener L1 (WR in pink and WOR in grey; left), and as labelled by the LENA software (WR in black and WOR in green; right). Raw and fitted probability distributions of step sizes from human labelled data and corresponding LENA labelled data are shown in panels (c1) through (f8). Infant data for step size distributions following a response (WR) are in panels c1–c4 and infant data for WOR are in panels c5–c8. Adult WR data are in panels d1–d4, and adult WOR data are in panels d5–d8. Fits for Infant WR data are in panels e1–e4 and infant WOR fits are in panels e5–e8. Adult WR fits are in panels f1–f4, and adult WOR fits are in panels f5–f8. The (f) in the legend indicates fitted as opposed to raw data (indicated by (d) in the legend). Note that both human-labelled infant and adult WR and WOR data/fits are given by pink and grey solid/dashed lines and are presented in the same subplots as their corresponding LENA-labelled data/fits, to allow for visual comparison of how well the curves do, or in a few cases do not, overlap.

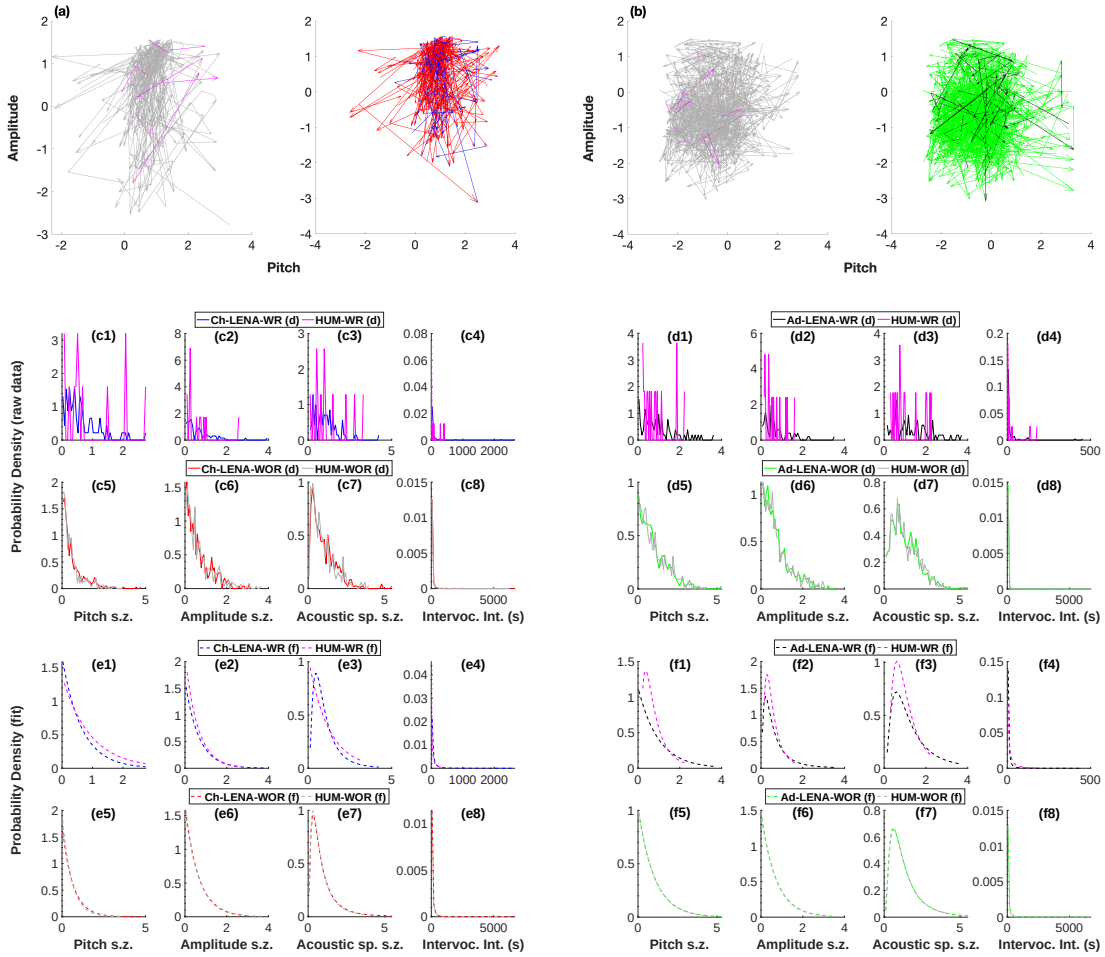

Figure B.16: **Human labelled data and corresponding LENA labelled data - Participant 340 at 183 days old; Listener 3.** (a) Acoustic space traversed by the infant as labelled by human listener L3 (WR in pink and WOR in grey; left), and as labelled by the LENA software (WR in blue and WOR in red; right). (b) Acoustic space traversed by the adult as labelled by human listener L3 (WR in pink and WOR in grey; left), and as labelled by the LENA software (WR in black and WOR in green; right). Raw and fitted probability distributions of step sizes from human labelled data and corresponding LENA labelled data are shown in panels (c1) through (f8). Infant data for step size distributions following a response (WR) are in panels c1–c4 and infant data for WOR are in panels c5–c8. Adult WR data are in panels d1–d4, and adult WOR data are in panels d5–d8. Fits for Infant WR data are in panels e1–e4 and infant WOR fits are in panels e5–e8. Adult WR fits are in panels f1–f4, and adult WOR fits are in panels f5–f8. The (f) in the legend indicates fitted as opposed to raw data (indicated by (d) in the legend). Note that both human-labelled infant and adult WR and WOR data/fits are given by pink and grey solid/dashed lines and are presented in the same subplots as their corresponding LENA-labelled data/fits, to allow for visual comparison of how well the curves do, or in a few cases do not, overlap.

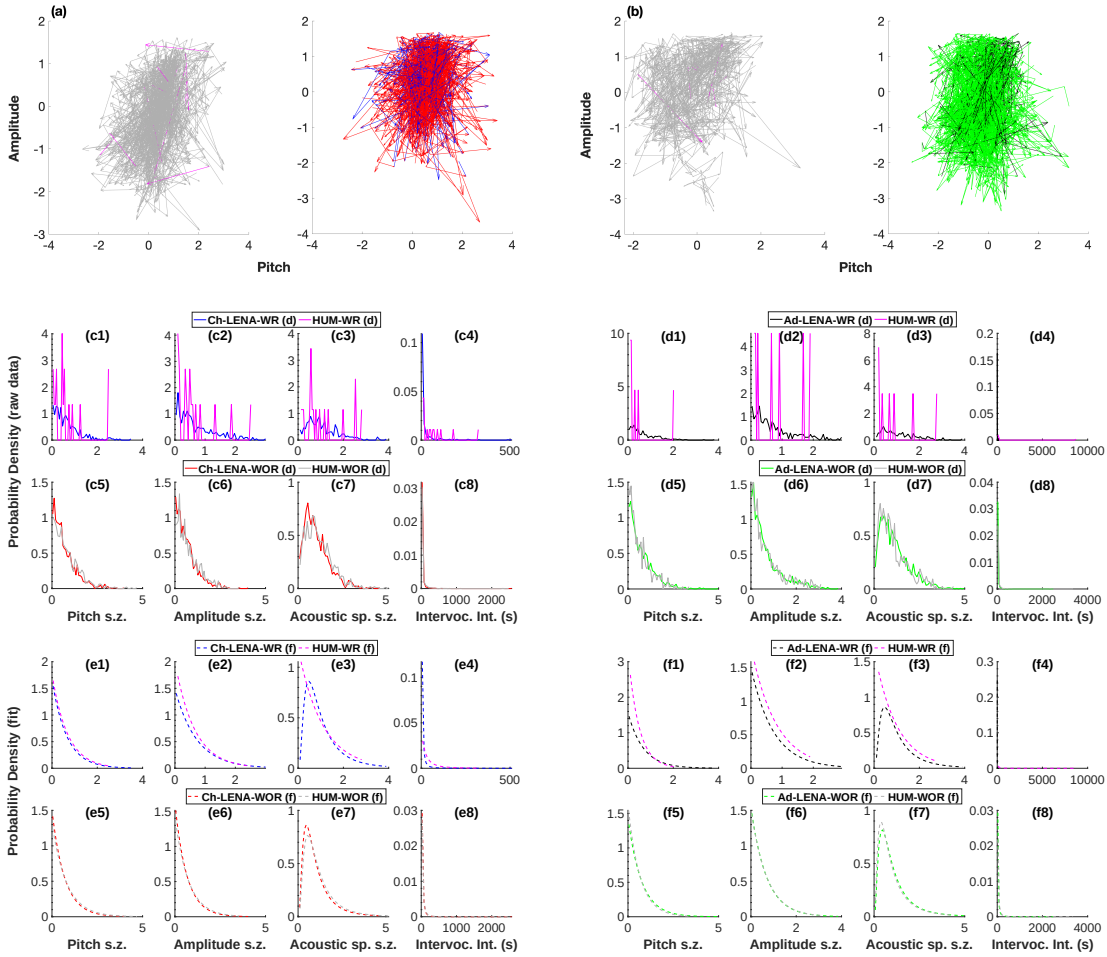

Figure B.17: **Human labelled data and corresponding LENA labelled data - Participant 530 at 95 days old; Listener 2.** (a) Acoustic space traversed by the infant as labelled by human listener L2 (WR in pink and WOR in grey; left), and as labelled by the LENA software (WR in blue and WOR in red; right). (b) Acoustic space traversed by the adult as labelled by human listener L2 (WR in pink and WOR in grey; left), and as labelled by the LENA software (WR in black and WOR in green; right). Raw and fitted probability distributions of step sizes from human labelled data and corresponding LENA labelled data are shown in panels (c1) through (f8). Infant data for step size distributions following a response (WR) are in panels c1–c4 and infant data for WOR are in panels c5–c8. Adult WR data are in panels d1–d4, and adult WOR data are in panels d5–d8. Fits for Infant WR data are in panels e1–e4 and infant WOR fits are in panels e5–e8. Adult WR fits are in panels f1–f4, and adult WOR fits are in panels f5–f8. The (f) in the legend indicates fitted as opposed to raw data (indicated by (d) in the legend). Note that both human-labelled infant and adult WR and WOR data/fits are given by pink and grey solid/dashed lines and are presented in the same subplots as their corresponding LENA-labelled data/fits, to allow for visual comparison of how well the curves do, or in a few cases do not, overlap.

B.5 What vocalisation acoustics and changes in vocalisation acoustics predict responses?

Table B.17: **Which vocalisation patterns received responses: results of statistical analysis with patterns of change in acoustics included.** βs are given with p-values in brackets. Significant results (at a significance level of 0.05) are in bold. The 'step' variables are the step sizes from the preceding vocalisation to the vocalisation in question. Note that acoustic step sizes are non-directional and may represent either increasing or decreasing amplitude or pitch. Infant ID was a random effect in all models. All values reported have been rounded to two decimal points wherever possible.

	Probability of infant receiving adult response	**Probability of adult receiving infant response**
Infant age	**-0.16**	-0.01
	(p<0.001)	(p=0.08)
Pitch	**-0.32**	**0.24**
	(p<0.001)	**(p<0.001)**
Amplitude	**0.18**	**0.44**
	(p<0.001)	**(p<0.001)**
Pitch step	**-0.02**	**-0.07**
	(p=0.002)	**(p<0.001)**
Amplitude step	**0.03**	**0.09**
	(p<0.001)	**(p<0.001)**
Time step	**-0.07**	**-0.03**
	(p<0.001)	**(p=0.02)**

Appendix C

Supplementary Information: The Fitness Trade-offs of Predation: When to Scavenge and When to Steal

C.1 Allometric relationships used in the SDP model

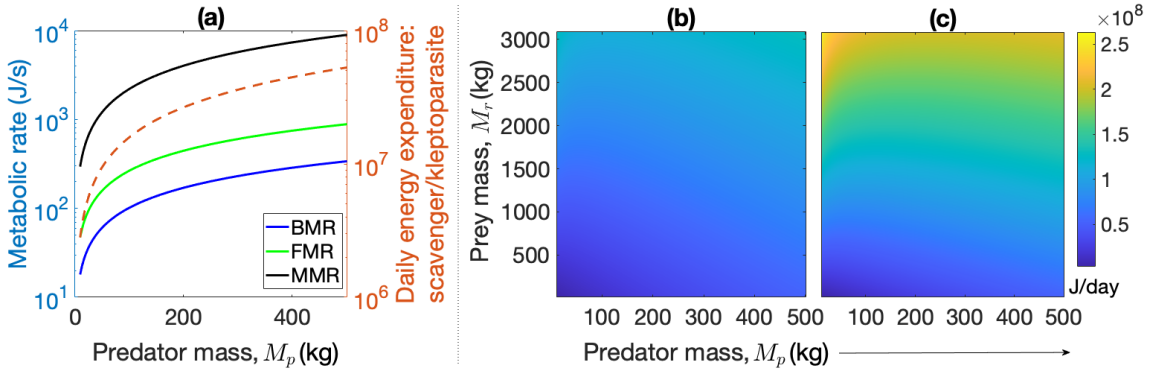

Figure C.1: **Allometric scaling of predator energy expenditure for different strategies**. (a) The left Y axis (blue) summarises how a predator's basal metabolic rate (BMR; blue solid line), field metabolic rate (FMR; green solid line), and maximal metabolic rate (MMR; black solid line) scales with predator body size. All metabolic rates are expressed in J/s. The right Y axis (red) shows how a scavenging or kleptoparasitic predator's energy expenditure (in J) over the course of a day scales with predator body mass (red dashed line). (b) shows a heat map of the energy expenditure associated with hunting over the course of a day (see colour bar; units of J/day) when the number of successful encounters with prey, $n^h = 1$, as a function of predator and prey mass. Here, the superscript h indicated hunting. (c) shows a similar heat map for $n^h = 2$.

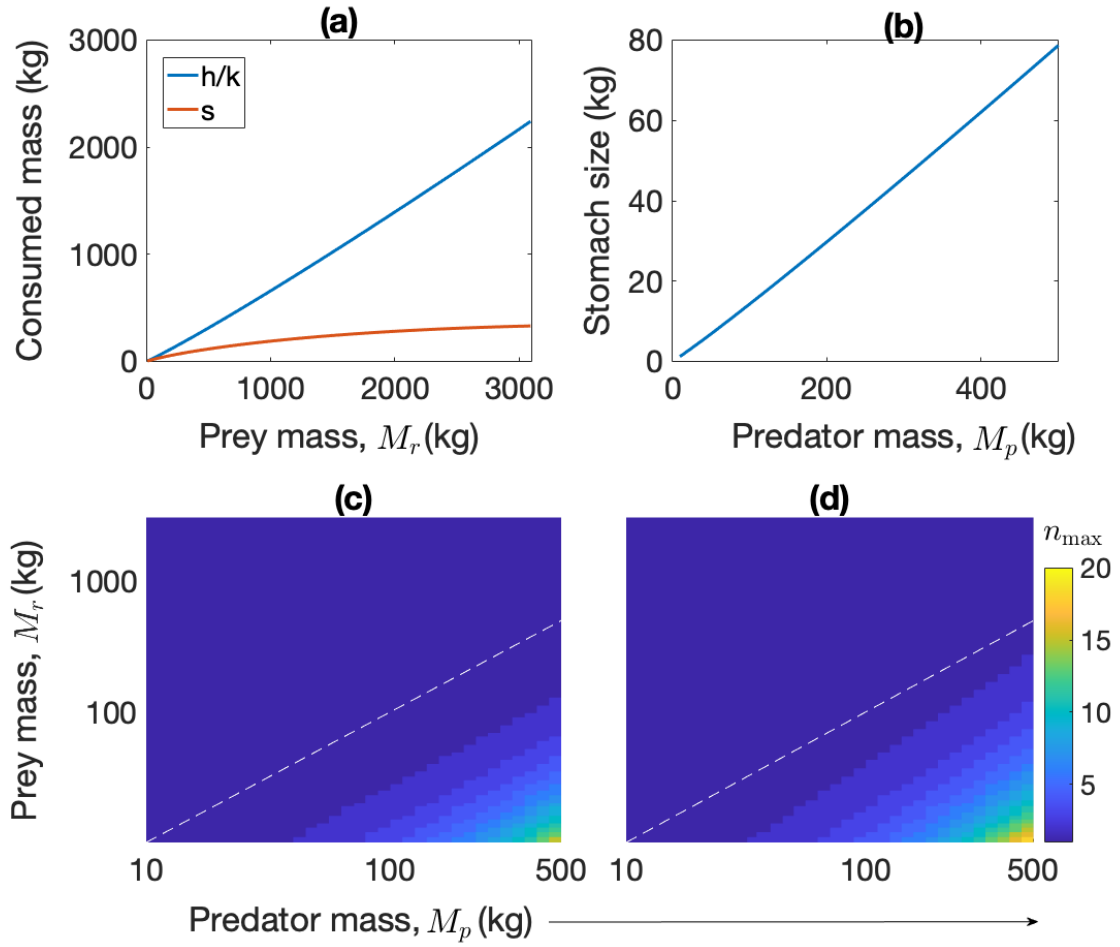

Figure C.2: **Allometric scaling of factors that affect predator energetic gains.** (a) Biomass consumed by a scavenger (red), and a hunter or kleptoparasite (blue) in one successful encounter is shown as a function of prey mass. The former is given by the sum of the prey's fat and muscle mass, while the latter is computed by subtracting the prey's fat, muscle, and skeletal mass from the prey mass. (b) Predator stomach size as a function of predator body size. (c) The maximum possible number of successful encounters during a 12 hour foraging bout, $n_{\max}$ (see colour bar) for hunting/stealing for all predator-prey mass pairings is shown as a heat map, as a function of predator mass, M_p, and prey mass, M_r. The white dashed line is the one-to-one mass line. (d) shows a similar heat map for $n_{\max}$ for scavenging. $n_{\max}$ is 1, regardless of strategy, when prey is larger than the predator (for the range of predator and prey masses investigated,). When prey is smaller than the predator, $n_{\max}$ is higher for scavengers when compared to the other two strategies. Closer to the one-to-one mass line, however, this number approaches 1 for all strategies.

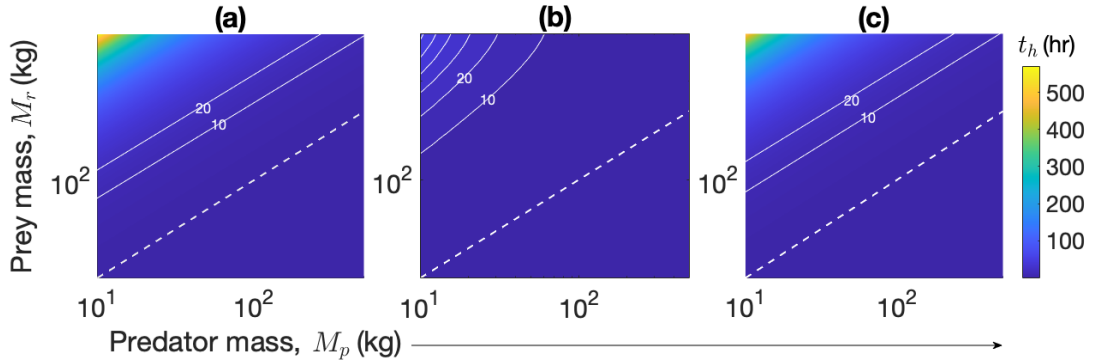

Figure C.3: **Allometric scaling of handling time for different strategies**. (a) The handling time, t_h (in hours), for a hunter as a function of predator and prey mass is shown as a heat map (see colour bar). The dashed white line is the one-to-one mass line. The contours for $t_h = 10, 20$ hours are shown. (b) shows a similar heat map for the handling time, t_h, for a scavenger. (c) shows a similar heat map for t_h for a kleptoparasite. When prey is smaller than the predator, t_h for hunting and stealing is less than two hours, while that for scavenging is less than an hour. In addition, t_h for stealing and hunting are comparable to each other while that for scavenging is much lower.

For hunting, the handling time is computed as the sum of the time it takes for the predator to pursue and subdue prey (estimated as 20% of the handling time according to the expression $t_h = 8912 M_p^{1.02} M_r$) and the time it takes to consume prey fat and muscle mass (estimated as 80% of $t_h = 8912 M_p^{1.02} r^h$), where r^h is the combined prey fat and muscle mass from one successful prey encounter).

For scavenging, the handling time is estimated as 80% of $t_h = 8912 M_p^{1.02} r^s$, where r^s is the mass of leftovers available to the scavenger in one successful encounter. This acocunts for the fact that the scavenger does not spend time pursuing and subduing prey.

Finally, in the case of kleptoparasitism, the handling time is computed as the sum of 10% of $t_h = 8912 M_p^{1.02} M_r$ (to account for the time spent stealing from the competitor, because the actual act of stealing is assumed to be a quick one, due to the high per unit time mortality associated with it) and 80% of $t_h = 8912 M_p^{1.02} r^k$, where $r^k = r^h$.

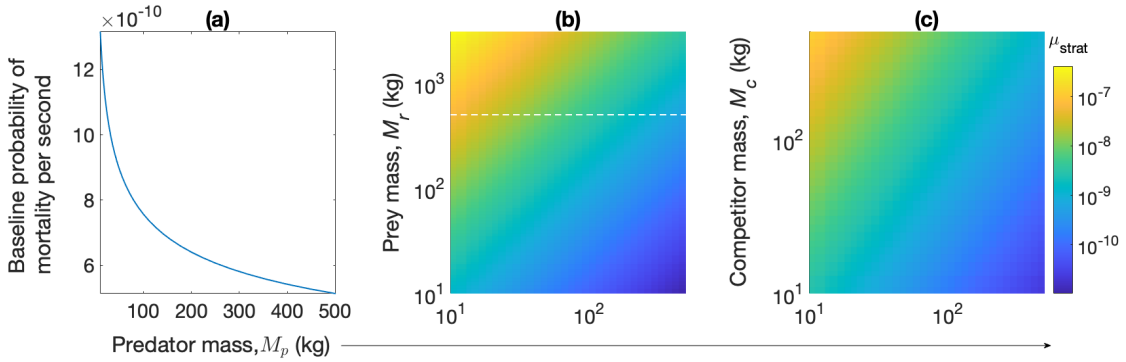

Figure C.4: **Allometric scaling of predator mortality for different foraging strategies.** (a) The probability of mortality per second for a scavenger as a function of predator body mass is shown. This serves as the baseline mortality that all predators incur at all times. During antagonistic encounters with a prey (or competitor) while hunting (or stealing), predators incur additional mortality risks. (b) Additional probability of mortality per second as incurred by predators during successful encounters with prey while hunting is shown as a function of predator and prey mass. The dashed horizontal line indicates the maximum competitor mass investigated in our model, to provide a comparison between the additional mortality due to hunting and the additional mortality due to stealing. (c) shows the additional mortality per second incurred during a successful encounter with the competitor as a function of predator mass and competitor mass. Note that the additional mortality due to stealing at competitor masses is comparable to that due to hunting at prey masses that are an order of magnitude higher.

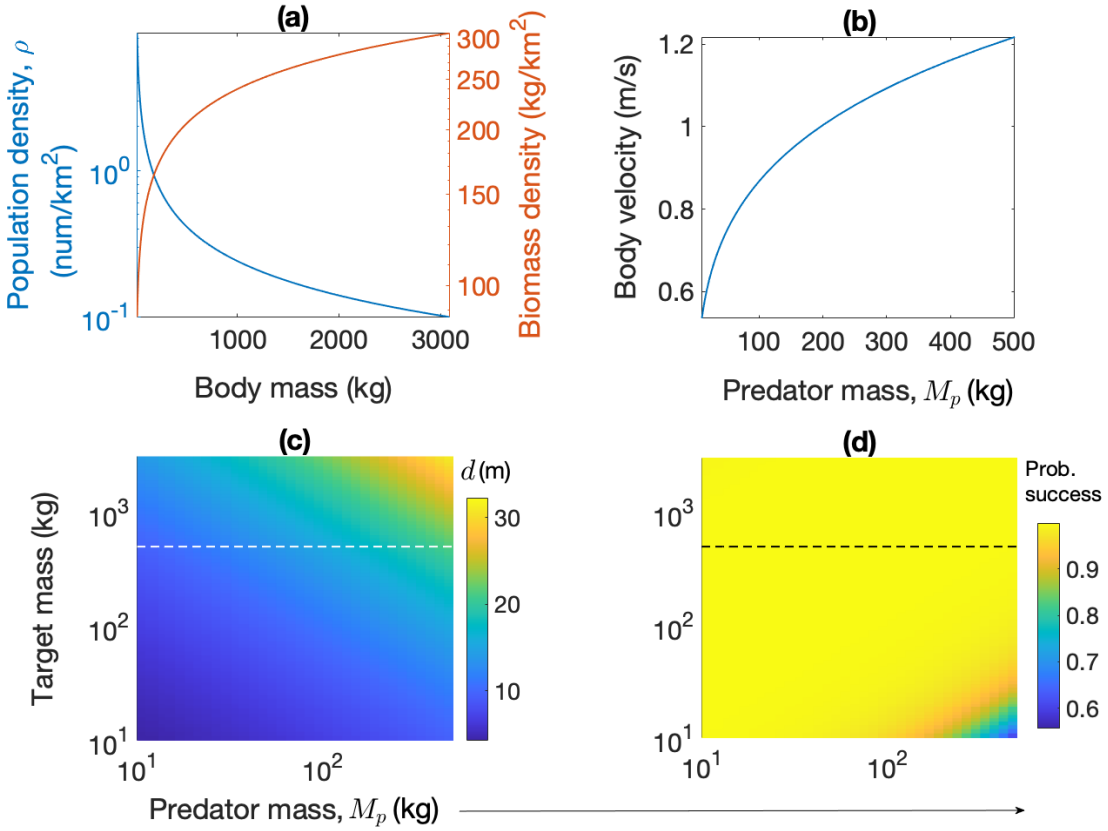

Figure C.5: **A summary of allometric scaling relationships used to compute encounter probability distributions for different strategies.** (a) The population density (in number of individuals per km^2) and the biomass density are plotted as a function of organismal body mass on the left and right Y axes, respectively. These relationships apply to predator, prey and competitor densities in our model. (b) The body velocity, v of the predator during a foraging bout is plotted as a function of predator mass. (c) The reaction distance, d of the predator for a target of given mass is depicted as a heat map. This relationship applies to both prey and competitor in our model, with the former being relevant to hunting, and the latter being relevant to scavenging and stealing. The dashed horizontal line indicates the maximum competitor mass investigated in this study. (d) shows the probability of a predator successfully subduing prey (in the case of hunting) or successfully stealing prey from a competitor (in the case of stealing) as a heat map. The dashed horizontal line indicates the maximum competitor mass investigated in this study.

C.2 Computing proportion of strategies

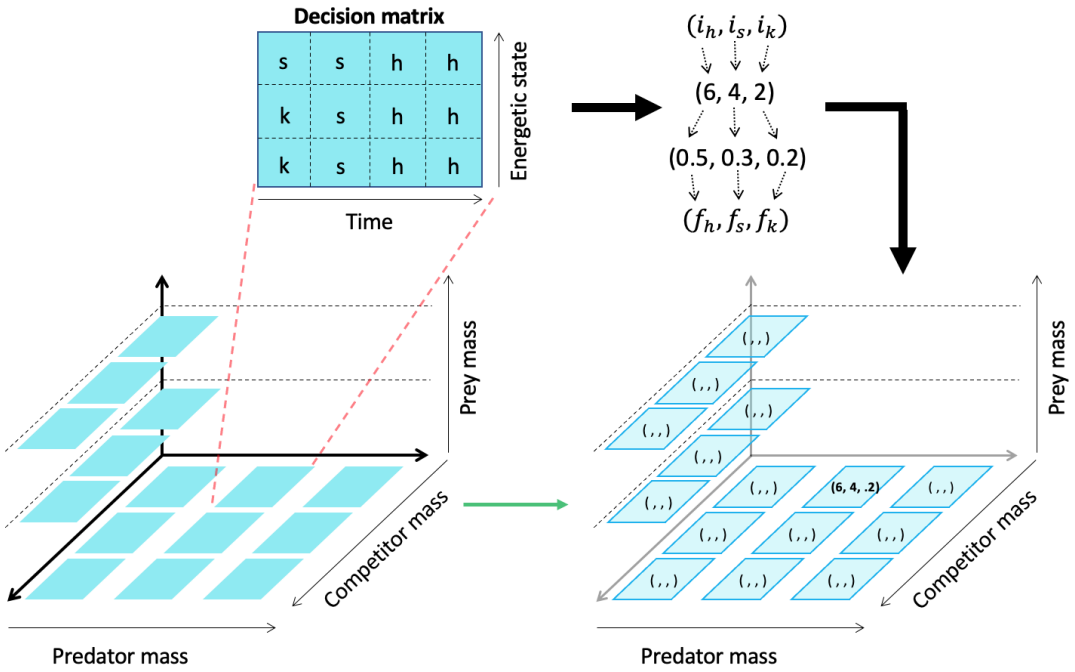

Figure C.6: **Schematic describing how the proportion of different foraging strategies is computed.** One trial of the simulation involves computing decision matrices for every combination of predator, prey, and competitor mass. If N_p, N_r, and N_v are the number of predator, prey, and competitor masses considered in the simulation, then there are $N_p \times N_r \times N_v$ decision matrices. Each decision matrix contains the strategy employed by the predator for discretised values of the state variables, energetic content and time (sample decision matrix shown). Here, h represents hunting, s represents scavenging, and k represents stealing (kleptoparasitism). Each decision matrix can be summarised using the numbers f_h, f_s, and f_k, which represent the fraction of hunting, scavenging, and kleptoparasitic 'decisions' in the matrix, respectively. Thus, the information in $N_p \times N_r \times N_v$ decision matrices can be reduced to $N_p \times N_r \times N_v$ vectors of the form $(f_h,\ f_s,\ f_k)$. Alternatively, the vectors $(i_h,\ i_s,\ i_k)$ can also be used, where i_h, i_s, and i_k are the number of hunting, scavenging, and stealing 'decisions' in the decision matrix.

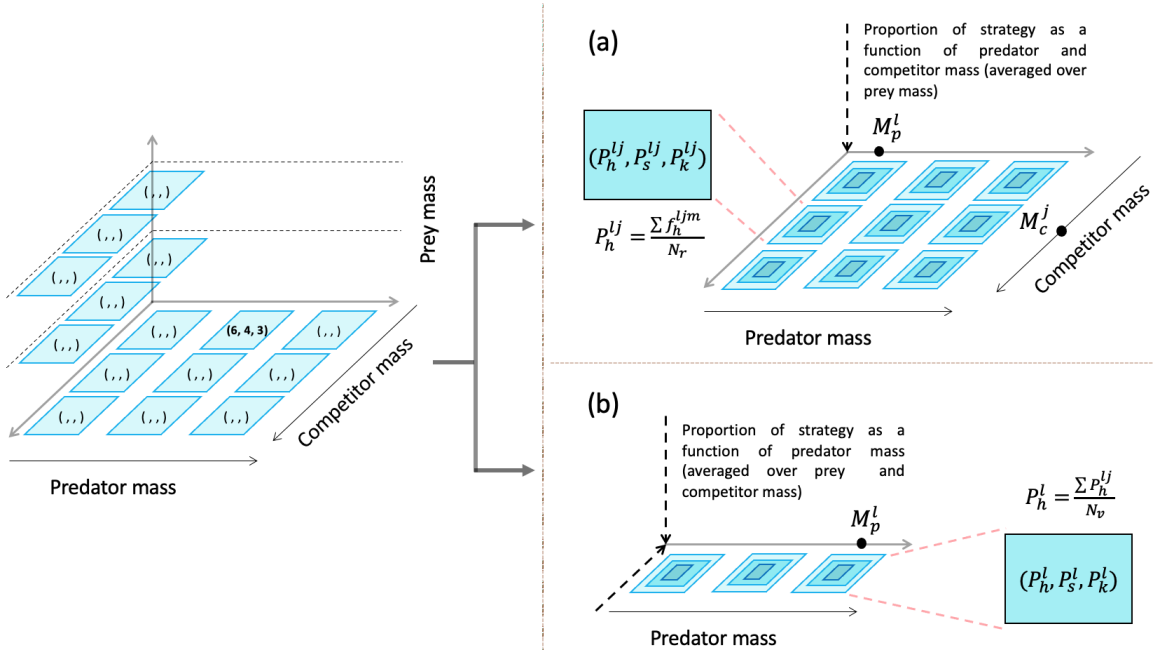

Figure C.7: **Schematic describing how the proportion of different foraging strategies is computed (contd.).** The vectors $(f_h^{ljm}, f_s^{ljm}, f_k^{ljm})$ can be used to compute the proportion of hunting (P_h), scavenging (P_s), and stealing (P_k) in different contexts. Here, the indices l, j, and m correspond to the predator, competitor, and prey mass, respectively. (a) To compute the proportion of a strategy as a function of predator and competitor mass, the vectors $(f_h^{ljm}, f_s^{ljm}, f_k^{ljm})$ are summed along the prey mass axis, as shown. Then, for predator mass M_p^l and competitor mass M_v^j, the proportion of hunting, P_h^{lj} is given by $\sum_m f_h^{ljm}/N_r$. Dividing by N_r normalises the summation. Proportions for the other two strategies can be computed similarly. (b) To compute the proportion of a strategy as a function of predator mass, the vectors $(f_h^{ljm}, f_s^{ljm}, f_k^{ljm})$ are summed along the prey mass and competitor mass axes. Then, for predator mass M_p^l the proportion of hunting, P_h^l is given by $\sum_{j,m} f_h^{ljm}/N_r N_v = \sum_j L_h^{ij}/N_v$. Dividing by $N_r N_v$ normalises the summation. Proportions for the other two strategies can be computed similarly.

C.3 More visualisations of average proportions of strategy

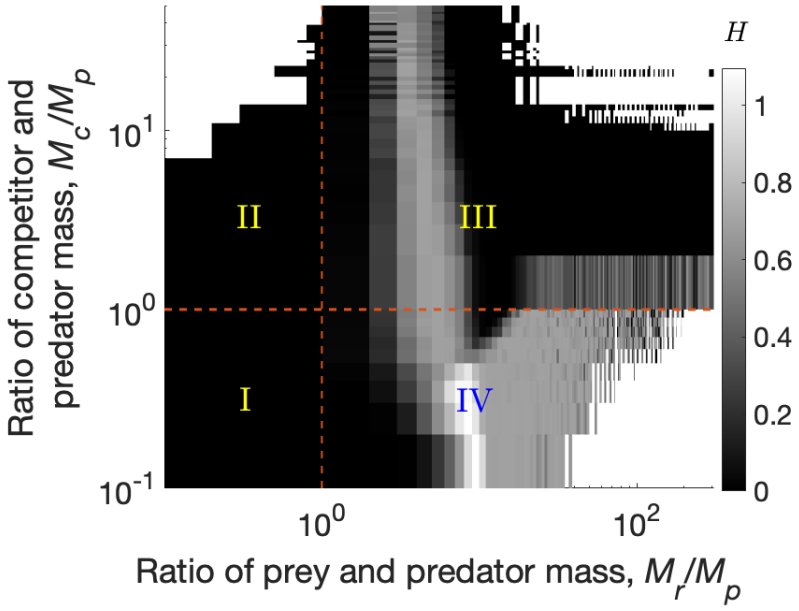

Figure C.8: **Relative predator size as a predictor of strategy diversity** Shannon's entropy, H of predator foraging strategy as a function of prey size and competitor size with respect to predator size. The vertical red dashed line indicates the line along which prey mass and predator mass are the same, while the horizontal red line is the line competitor mass equal to predator mass. Strategy diversity is highest in quadrant IV, when competitors are smaller than or about the same size as the predator, and prey is about 10 times larger than the predator.

Bibliography

[1] VPS Ritwika, Gina M Pretzer, Sara Mendoza, Christopher Shedd, Christopher T Kello, Ajay Gopinathan, and Anne S Warlaumont. Exploratory dynamics of vocal foraging during infant-caregiver communication. *Scientific reports*, 10(1):1–14, 2020.

[2] Alexandra-Maria Klein, Bernard E Vaissiere, James H Cane, Ingolf Steffan-Dewenter, Saul A Cunningham, Claire Kremen, and Teja Tscharntke. Importance of pollinators in changing landscapes for world crops. *Proceedings of the royal society B: biological sciences*, 274(1608):303–313, 2007.

[3] Marcelo A Aizen, Lucas A Garibaldi, Saul A Cunningham, and Alexandra M Klein. How much does agriculture depend on pollinators? lessons from long-term trends in crop production. *Annals of botany*, 103(9):1579–1588, 2009.

[4] Jim Vafidis, Jeremy Smith, and Robert Thomas. Climate change and insectivore ecology. *eLS*, pages 1–8, 2019.

[5] Davide Normanno, Lydia Boudarène, Claire Dugast-Darzacq, Jiji Chen, Christian Richter, Florence Proux, Olivier Bénichou, Raphaël Voituriez, Xavier Darzacq, and Maxime Dahan. Probing the target search of dna-binding proteins in mammalian cells using tetr as model searcher. *Nature communications*, 6(1):1–10, 2015.

[6] Charles W Fox, Derek A Roff, and Daphne J Fairbairn. *Evolutionary ecology: concepts and case studies*. Oxford University Press, 2001.

[7] Gandhimohan M Viswanathan, Marcos GE Da Luz, Ernesto P Raposo, and H Eugene Stanley. *The physics of foraging: an introduction to random searches and biological encounters*. Cambridge University Press, 2011.

[8] Frederic Bartumeus and Jordi Catalan. Optimal search behavior and classic foraging theory. *Journal of Physics A: Mathematical and Theoretical*, 42(43):434002, 2009.

[9] Gandhimohan M Viswanathan, V Afanasyev, Sergey V Buldyrev, EJ Murphy, PA Prince, and H Eugene Stanley. Lévy flight search patterns of wandering albatrosses. *Nature*, 381(6581):413–415, 1996.

[10] Gandimohan M Viswanathan, Sergey V Buldyrev, Shlomo Havlin, MGE Da Luz, EP Raposo, and H Eugene Stanley. Optimizing the success of random searches. *nature*, 401(6756):911–914, 1999.

[11] Gandhimohan M Viswanathan, EP Raposo, and MGE Da Luz. Lévy flights and superdiffusion in the context of biological encounters and random searches. *Physics of Life Reviews*, 5(3):133–150, 2008.

[12] Roman Stocker and Justin R Seymour. Ecology and physics of bacterial chemotaxis in the ocean. *Microbiology and Molecular Biology Reviews*, 76(4):792–812, 2012.

[13] Bryan E Kerster, Theo Rhodes, and Christopher T Kello. Spatial memory in foraging games. *Cognition*, 148:85–96, 2016.

[14] Thomas T Hills, Michael N Jones, and Peter M Todd. Optimal foraging in semantic memory. *Psychological review*, 119(2):431, 2012.

[15] Theo Rhodes and Michael T Turvey. Human memory retrieval as lévy foraging. *Physica A: Statistical Mechanics and its Applications*, 385(1):255–260, 2007.

[16] Gijs Koot, Mirjam AA Huis in Veld, Joost Hendricksen, Rianne Kaptein, Arnout De Vries, and Egon L Van Den Broek. Foraging online social networks. In *2014 IEEE Joint Intelligence and Security Informatics Conference*, pages 312–315. IEEE, 2014.

[17] Alan Wexelblat and Pattie Maes. Footprints: History-rich tools for information foraging. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 270–277, 1999.

[18] MW Hayward, W Jedrzejewski, and B Jedrzejewska. Prey preferences of the tiger p anthera tigris. *Journal of Zoology*, 286(3):221–231, 2012.

[19] MW Hayward, Philippe Henschel, John O'Brien, Markus Hofmeyr, Guy Balme, and Graham IH Kerley. Prey preferences of the leopard (panthera pardus). *Journal of Zoology*, 270(2):298–313, 2006.

[20] Reuven Dukas and Leslie A Real. Learning foraging tasks by bees: a comparison between social and solitary species. *Animal Behaviour*, 42(2):269–276, 1991.

[21] James S Diana. Diel activity pattern and swimming speeds of northern pike (esox lucius) in lac ste. anne, alberta. *Canadian Journal of Fisheries and Aquatic Sciences*, 37(9):1454–1458, 1980.

[22] Jean-Marc Thiollay. Comparative foraging behavior between solitary and flocking insectivores in a neotropical forest: does vulnerability matter. *Ornitologia neotropical*, 14:47–65, 2003.

[23] David W Stephens, Joel S Brown, and Ronald C Ydenberg. *Foraging: behavior and ecology*. University of Chicago Press, 2008.

[24] Robert H MacArthur and Eric R Pianka. On optimal use of a patchy environment. *The American Naturalist*, 100(916):603–609, 1966.

[25] J Merritt Emlen. The role of time and energy in food preference. *The American Naturalist*, 100(916):611–617, 1966.

[26] Eric R Pianka. Convexity, desert lizards, and spatial heterogeneity. *Ecology*, 47(6):1055–1059, 1966.

[27] Gad Perry and Eric R Pianka. Animal foraging: past, present and future. *Trends in Ecology & Evolution*, 12(9):360–364, 1997.

[28] Graham H Pyke, H Ronald Pulliam, and Eric L Charnov. Optimal foraging: a selective review of theory and tests. *The quarterly review of biology*, 52(2):137–154, 1977.

[29] Graham H Pyke. Optimal foraging theory: a critical review. *Annual review of ecology and systematics*, 15(1):523–575, 1984.

[30] Marc Mangel and Colin W Clark. Towards a unifield foraging theory. *Ecology*, 67(5):1127–1138, 1986.

[31] Luc-Alain Giraldeau and Guy Beauchamp. Food exploitation: searching for the optimal joining policy. *Trends in Ecology & Evolution*, 14(3):102–106, 1999.

[32] Kunal Bhattacharya and Tamás Vicsek. To join or not to join: collective foraging strategies. In *Journal of Physics: conference series*, volume 638, page 012015. IOP Publishing, 2015.

[33] Kunal Bhattacharya and Tamás Vicsek. Collective foraging in heterogeneous landscapes. *Journal of The Royal Society Interface*, 11(100):20140674, 2014.

[34] MW Hayward. Prey preferences of the spotted hyaena (crocuta crocuta) and degree of dietary overlap with the lion (panthera leo). *Journal of Zoology*, 270(4):606–614, 2006.

[35] Tim Caro. *Cheetahs of the Serengeti Plains: group living in an asocial species*. University of Chicago Press, 1994.

[36] Jacob D Davidson and Deborah M Gordon. Spatial organization and interactions of harvester ants during foraging activity. *Journal of The Royal Society Interface*, 14(135):20170413, 2017.

[37] WG Wilson, CW Osenberg, RJ Schmitt, and RM Nisbet. Complementary foraging behaviors allow coexistence of two consumers. *Ecology*, 80(7):2358–2372, 1999.

[38] William L Vickery, Luc-Alain Giraldeau, Jennifer J Templeton, Donald L Kramer, and Colin A Chapman. Producers, scroungers, and group foraging. *The american naturalist*, 137(6):847–863, 1991.

[39] Tristan Ursell, Rosanna Man Wah Chau, Susanne Wisen, Devaki Bhaya, and Kerwyn Casey Huang. Motility enhancement through surface modification is sufficient for cyanobacterial community organization during phototaxis. *PLoS computational biology*, 9(9):e1003205, 2013.

[40] V Elangovan, G Marimuthu, and TH Kunz. Temporal patterns of individual and group foraging behaviour in the short-nosed fruit bat, cynopterus sphinx, in south india. *Journal of Tropical Ecology*, 15(5):681–687, 1999.

[41] HJ De Knegt, GM Hengeveld, F Van Langevelde, WF De Boer, and KP Kirkman. Patch density determines movement patterns and foraging efficiency of large herbivores. *Behavioral Ecology*, 18(6):1065–1072, 2007.

[42] Jong-Soon Choi, Young-Ho Chung, Yoon-Jung Moon, Changhoon Kim, Masakatsu Watanabe, Pill-Soon Song, Cheol-O Joe, Lawrence Bogorad, and Young Mok Park. Photomovement of the gliding cyanobacterium synechocystis sp. pcc 6803. *Photochemistry and photobiology*, 70(1):95–102, 1999.

[43] C. Moulin-Frier, S. M. Nguyen, and P-Y. Oudeyer. Self-organization of early vocal development in infants and machines: The role of intrinsic motivation. *Frontiers in Psychology*, 4:1006, 2014.

[44] A. S. Warlaumont, G. Westermann, E. H. Buder, and D. K. Oller. Prespeech motor learning in a neural network using reinforcement. *Neural Networks*, 38:64–75, 2013.

[45] A. S. Warlaumont and M. K. Finnegan. Learning to produce syllabic speech sounds via reward-modulated neural plasticity. *PLOS ONE*, 11:e0145096, 2016.

[46] Y. Yoshikawa, M. Asada, K. Hosoda, and J. Koga. A constructivist approach to infants' vowel acquisition through mother-infant interaction. *Connection Science*, 15:245–258, 2003.

[47] K. Miura, Y. Yoshikawa, and M. Asada. Vowel acquisition based on an auto-mirroring bias with a less imitative caregiver. *Advanced Robotics*, 26:23–44, 2012.

[48] I. S. Howard and P. Messum. Learning to pronounce first words in three languages: An investigation of caregiver and infant behavior using a computational model of an infant. *PLOS ONE*, 9:e110334, 2014.

[49] B. J. Kröger, J. Kannampuzha, and E. Kaufmann. Associative learning and self-organization as basic principles for simulating speech acquisition, speech production, and speech perception. *EPJ Nonlinear Biomedical Physics*, 2:2, 2014.

[50] Gáspár Jékely, Julien Colombelli, Harald Hausen, Keren Guy, Ernst Stelzer, François Nédélec, and Detlev Arendt. Mechanism of phototaxis in marine zooplankton. *Nature*, 456(7220):395–399, 2008.

[51] Gáspár Jékely. Evolution of phototaxis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1531):2795–2808, 2009.

[52] Nortin M Hadler. Heritability and phototaxis in drosophila melanogaster. *Genetics*, 50(6):1269, 1964.

[53] Masahiko Ikeuchi and Satoshi Tabata. Synechocystis sp. pcc 6803—a useful tool in the study of the genetics of cyanobacteria. *Photosynthesis research*, 70(1):73–83, 2001.

[54] Henning Knoop, Yvonne Zilliges, Wolfgang Lockau, and Ralf Steuer. The metabolic network of synechocystis sp. pcc 6803: systemic properties of autotrophic growth. *Plant physiology*, 154(1):410–422, 2010.

[55] Matthew Burriesci and Devaki Bhaya. Tracking phototactic responses and modeling motility of synechocystis sp. strain pcc6803. *Journal of Photochemistry and Photobiology B: Biology*, 91(2-3):77–86, 2008.

[56] Rosanna Man Wah Chau, Tristan Ursell, Shuo Wang, Kerwyn Casey Huang, and Devaki Bhaya. Maintenance of motility bias during cyanobacterial phototaxis. *Biophysical journal*, 108(7):1623–1632, 2015.

[57] Conrad W Mullineaux. How do cyanobacteria sense and respond to light? *Molecular microbiology*, 41(5):965–971, 2001.

[58] Devaki Bhaya. Light matters: phototaxis and signal transduction in unicellular cyanobacteria. *Molecular microbiology*, 53(3):745–754, 2004.

[59] Shizue Yoshihara and Masahiko Ikeuchi. Phototactic motility in the unicellular cyanobacterium synechocystis sp. pcc 6803. *Photochemical & Photobiological Sciences*, 3(6):512–518, 2004.

[60] Devaki Bhaya, Akiko Takahashi, Payam Shahi, and Arthur R Grossman. Novel motility mutants of synechocystis strain pcc 6803 generated by in vitro transposon mutagenesis. *Journal of Bacteriology*, 183(20):6140–6143, 2001.

[61] Brita Fiedler, Thomas Börner, and Annegret Wilde. Phototaxis in the cyanobacterium synechocystis sp. pcc 6803: role of different photoreceptors. *Photochemistry and photobiology*, 81(6):1481–1488, 2005.

[62] Ji-Young Song, Hye Sun Cho, Jung-Il Cho, Jong-Seong Jeon, J Clark Lagarias, and Youn-Il Park. Near-uv cyanobacteriochrome signaling system elicits negative phototaxis in the cyanobacterium synechocystis sp. pcc 6803. *Proceedings of the National Academy of Sciences*, 108(26):10780–10785, 2011.

[63] Devaki Bhaya, Nicole R Bianco, Donald Bryant, and Arthur Grossman. Type iv pilus biogenesis and motility in the cyanobacterium synechocystis sp. pcc6803. *Molecular microbiology*, 37(4):941–951, 2000.

[64] Wing-On Ng, Arthur R Grossman, and Devaki Bhaya. Multiple light inputs control phototaxis in synechocystis sp. strain pcc6803. *Journal of Bacteriology*, 185(5):1599–1607, 2003.

[65] Berrin Tansel and Derya Z Tansel. Adhesion strength and spreading characteristics of eps on membrane surfaces during lateral and central growth. *Colloids and Surfaces B: Biointerfaces*, 111:594–599, 2013.

[66] K Becker. Exopolysaccharide production and attachment strength of bacteria and diatoms on substrates with different surface tensions. *Microbial ecology*, 32(1):23–33, 1996.

[67] Altan Ozkan and Halil Berberoglu. Physico-chemical surface properties of microalgae. *Colloids and surfaces B: Biointerfaces*, 112:287–293, 2013.

[68] Sara Pereira, Andrea Zille, Ernesto Micheletti, Pedro Moradas-Ferreira, Roberto De Philippis, and Paula Tamagnini. Complexity of cyanobacterial exopolysaccharides: composition, structures, inducing factors and putative genes involved in their biosynthesis and assembly. *FEMS microbiology reviews*, 33(5):917–941, 2009.

[69] Jing Li, Daoyan Fang, Rumeng Ye, Changfang Zhou, and Pengfu Li. The released polysaccharide inhibits cell aggregation and biofilm formation in the cyanobacterium synechocystis sp. pcc 6803. *European Journal of Phycology*, 56(2):119–128, 2021.

[70] D. K. Oller. The emergence of the sounds of speech in infancy. In G. H. Yeni-Komshian, J. F. Kavanagh, and C. A. Ferguson, editors, *Child Phonology*, volume 1: Production, pages 93–112. Academic Press, New York, 1980.

[71] D. K. Oller. *The emergence of the speech capacity*. Lawrence Erlbaum Associates, Mahwah, 2000.

[72] R. E. Stark. Stages of speech development in the first year of life. In G. H. Yeni-Komshian, J. F. Kavanagh, and C. A. Ferguson, editors, *Child Phonology*, volume 1: Production, pages 73–92. Academic Press, New York, 1980.

[73] F. J. Koopmans-van Beinum and J. M. van der Stelt. Early stages in the development of speech movements. In B. Lindblom and R. Zetterström, editors, *Precursors of early speech*, pages 37–50. Stockton Press, New York, 1986.

[74] E. H. Buder, A. S. Warlaumont, and D. K. Oller. An acoustic phonetic catalog of prespeech vocalizations from a developmental perspective. In B. Peter and A. N. MacLeod, editors, *Comprehensive perspectives on child speech development and disorders: Pathways from linguistic theory to clinical practice*, pages 103–134. Nova Science Publishers, New York, 2013.

[75] J. L. Locke. Babbling and early speech: Continuity and individual differences. *First Language*, 9:191–206, 1989.

[76] M. A. Vihman, M. A. Macken, R. Miller, H. Simmons, and J. Miller. From babbling to speech: A re-assessment of the continuity issue. *Language*, 61:397–445, 1985.

[77] L. Ménard, J. L. Schwartz, and L. J. Boë. Role of vocal tract morphology in speech development: Perceptual targets and sensorimotor maps for synthesized french vowels from birth to adulthood. *Journal of Speech, Language, and Hearing Research*, 47:1059–1080, 2004.

[78] A. S. Warlaumont. Modeling the emergence of syllabic structure. *Journal of Phonetics*, 53:61–65, 2015.

[79] M. H. Goldstein, A. P. King, and M. J. West. Social interaction shapes babbling: Testing parallels between birdsong and speech. *Proceedings of the National Academy of Sciences of the United States of America*, 100:8030–8035, 2003.

[80] S. Nathani and R. E. Stark. Can conditioning procedures yield representative infant vocalizations in the laboratory? *First Language*, 16:365–387, 1996.

[81] K. Bloom. Quality of adult vocalizations affects the quality of infant vocalizations. *Journal of Child Language*, 15:469–480, 1988.

[82] M. H. Goldstein and J. A. Schwade. Social feedback to infants' babbling facilitates rapid phonological learning. *Psychological Science*, 19:515–523, 2008.

[83] A. S. Warlaumont, J. A. Richards, J. Gilkerson, and D. K. Oller. Social feedback to infants' babbling facilitates rapid phonological learning. *Psychological Science*, 19:515–523, 2008.

[84] K. Bloom, A. Russell, and K. Wassenberg. Turn taking affects the quality of infant vocalizations. *Journal of Child Language*, 14:211–227, 1987.

[85] C. S. Tamis-LeMonda, M. H. Bornstein, and L. Baumwell. Maternal responsiveness and children's achievement of language milestones. *Child Development*, 72:748–767, 2001.

[86] J. Gros-Louis, M. J. West, and A. P. King. Maternal responsiveness and the development of directed vocalizing in social interactions. *Infancy*, 19:385–408, 2014.

[87] J. Jaffe, B. Beebe, S. Feldstein, C. L. Crown, and M. D. Jasnow. Rhythms of dialogue in infancy: Coordinated timing in development. *Monographs of the Society for Research in Child Development*, 66:vii–viii, 1–132, 2001.

[88] N. B. Leezenbaum, S. B. Campbell, D. Butler, and J. M. Iverson. Maternal verbal responses to communication of infants at low and heightened risk of autism. *Autism*, 18:694–703, 2013.

[89] S. L. Velleman, L. Mangipudi, and J. L. Locke. Prelinguistic phonetic contingency: Data from down syndrome. *First Language*, 9:150–173, 1989.

[90] D. K. Oller, E. H. Buder, H. L. Ramsdell, A. S. Warlaumont, L. Chorna, and R. Bakeman. Functional flexibility of infant vocalization and the emergence of language. *Proceedings of the National Academy of Sciences of the United States of America*, 10:6318–6323, 2013.

[91] D. H. Abney, A. S. Warlaumont, D. K. Oller, S. Wallot, and C. T. Kello. Multiple coordination patterns in infant and adult vocalizations. *Infancy*, 22:514–539, 2017.

[92] S. Falk and C. T. Kello. Hierarchical organization in the temporal structure of infant-direct speech and song. *Cognition*, 163:80–86, 2017.

[93] G. S. Katz, J. F. Cohn, and C. A. Moore. A combination of vocal f0 dynamic and summary features discriminates between three pragmatic categories of infant-directed speech. *Child Development*, 67:205–217, 1996.

[94] A. Fernald, T. Taeschner, J. Dunn, M. Papousek, and B. de Boysson-Bardies. A cross-language study of prosodic modifications in mothers' and fathers' speech to preverbal infants. *Journal of Child Language*, 16:477–501, 1989.

[95] E. A. Piazza, M. C. Iordan, and C. Lew-Williams. Mothers consistently alter their unique vocal fingerprints when communicating with infants. *Current Biology*, 27:3162–3167.e3, 2017.

[96] S. Schuster, S. Pancoast, M. Ganjoo, M. C. Frank, and D. Jurafsky. Speaker-independent detection of child-directed speech. In *IEEE Spoken Language Technology Workshop (SLT)*, pages 366–371, 2014.

[97] R. M. Golinkoff, D. D. Can, M. Soderstrom, and K. Hirsh-Pasek. (baby)talk to me: The social context of infant-directed speech and its effects on early language acquisition. *Current Directions in Psychological Science*, 24:339–344, 2015.

[98] P. De Palma and M. VanDam. Using automatic speech processing to analyze fundamental frequency of child-directed speech stored in a very large audio corpus. In *IEEE Proceedings of the Joint 17th World Congress of International Fuzzy Systems Association and 9th International Conference on Soft Computing and Intelligent Systems (IFSA-SCIS 2017)*, 2017.

[99] L. A. Van Egeren, M. S. Barratt, and M. A. Roach. Mother-infant responsiveness: Timing, mutual regulation, and interactional context. *Developmental Psychology*, 37:684–697, 2001.

[100] M. H. Bornstein, D. L. Putnick, L. R. Cote, O. M. Haynes, and J. T. D. Suwalsky. Mother-infant contingent vocalizations in 11 countries. *Psychological Science*, 26:1272–1284, 2015.

[101] G. M. Pretzer, L. D. Lopez, E. A. Walle, and A. S. Warlaumont. Infant-adult vocal interaction dynamics depend on infant vocal type, child-directedness of adult speech, and timeframe. *Infant Behavior and Development*, 57:101325, 2019.

[102] L. Shneidman, R. Todd, and A. Woodward. Why do child-directed interactions support imitative learning in young children? *PLOS ONE*, 9:e110891, 2014.

[103] T. T. Hills, M. N. Jones, and P. M. Todd. Optimal foraging in semantic memory. *Psychological Review*, 119:431–440, 2012.

[104] T. T. Hills, P. M. Todd, D. Lazer, A. D. Redish, and I. D. Couzin. Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, 19:46–54, 2015.

[105] P. Montez, G. Thompson, and C. T. Kello. The role of semantic clustering in optimal memory foraging. *Cognitive Science*, 39:1925–1939, 2015.

[106] T. Rhodes and M. T. Turvey. Human memory retrieval as lévy foraging. *Physica A: Statistical Mechanics and its Applications*, 385:255–260, 2007.

[107] F. Radicchi, A. Baronchelli, and L. A. N. Amaral. Rationality, irrationality and escalating behavior in lowest unique bid auctions. *PLOS ONE*, 7:e29910, 2012.

[108] G. M. Viswanathan, M. G. E. da Luz, E. P. Raposo, and E. H. Stanley. *The physics of foraging: An introduction to random searches and biological encounters*. Cambridge University Press, Cambridge, 2011.

[109] B. E. Kerster, T. Rhodes, and C. T. Kello. Spatial memory in foraging games. *Cognition*, 148:85–96, 2016.

[110] N. E. Humphries, N. Queiroz, J. R. M. Dyer, N. G. Pade, M. K. Musyl, K. M. Schaefer, D. W. Fuller, J. M. Brunnschweiler, T. K. Doyle, J. D. R. Houghton, G. C. Hays, C. S. Jones, L. R. Noble, V. J. Wearmouth, E. J. Southall, and D. W. Sims. Environmental context explains lévy and brownian movement patterns of marine predators. *Nature*, 465:1066–1069, 2010.

[111] J. A. Dixon, J. G. Holden, D. Mirman, and D. G. Stephen. Multifractal dynamics in the emergence of cognitive structure. *Topics in Cognitive Science*, 4:51–62, 2012.

[112] D. G. Stephen, D. Mirman, J. S. Magnuson, and J. A. Dixon. Lévy-like diffusion in eye movements during spoken-language comprehension. *Physical Review E: Statistical, Nonlinear, and Soft Matter Physics*, 79:056114, 2009.

[113] H. Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19:716–723, 1974.

[114] J. Gros-Louis and J. L. Miller. From 'ah' to 'bah': Social feedback loops for speech sounds at key points of developmental transition. *Journal of Child Language*, 45:807–825, 2018.

[115] A. Fernald. Four-month-old infants prefer to listen to motherese. *Infant Behavior and Development*, 8:181–195, 1985.

[116] E. Bergelson, M. Casillas, M. Soderstrom, A. Seidl, A. S. Warlaumont, and A. Amatuni. What do North American babies hear? a large-scale cross-corpus analysis. *Developmental Science*, 22:e12724, 2018.

[117] E. H. Buder, L. Chorna, D. K. Oller, and R. Robinson. Vibratory regime classification of infant phonation. *Journal of Voice*, 22:553–564, 2008.

[118] M. Papoušek and H. Papoušek. Forms and functions of vocal matching in interactions between mothers and their precanonical infants. *First Language*, 9:137–157, 1989.

[119] A. Fernald and P. Kuhl. Acoustic determinants of infant preference for motherese speech. *Infant Behavior and Development*, 10:279–293, 1987.

[120] D. Xu, U. Yapanel, and S. Gray. Reliability of the LENA$^{\text{TM}}$ language environment analysis system in young children's natural home environment. Technical report, LENA Foundation, Boulder, CO, 2009.

[121] P. Boersma and D. Weenink. Praat: doing phonetics by computer, 2018.

[122] Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. Fitting linear mixed-effects models using lme4. *arXiv preprint arXiv:1406.5823*, 2014.

[123] Shomen Mukherjee and Michael R Heithaus. Dangerous prey and daring predators: a review. *Biological Reviews*, 88(3):550–563, 2013.

[124] Chris Carbone, Nathalie Pettorelli, and Philip A Stephens. The bigger they come, the harder they fall: body size and prey abundance influence predator–prey ratios. *Biology letters*, 7(2):312–315, 2011.

[125] Simon Jennings and Steve Mackinson. Abundance–body mass relationships in size-structured food webs. *Ecology Letters*, 6(11):971–974, 2003.

[126] Chris Carbone, J Marcus Rowcliffe, Guy Cowlishaw, and Nick JB Isaac. The scaling of abundance in consumers and their resources: implications for the energy equivalence rule. *The American Naturalist*, 170(3):479–484, 2007.

[127] Uttam Bhat, Christopher P. Kempes, and Justin D. Yeakel. Scaling the risk landscape drives optimal life-history strategies and the evolution of grazing. *Proceedings of the National Academy of Sciences*, 2020.

[128] ARE Sinclair, Simon Mduma, and Justin S Brashares. Patterns of predation in a diverse predator–prey system. *Nature*, 425(6955):288–290, 2003.

[129] Chris Carbone, Amber Teacher, and J Marcus Rowcliffe. The costs of carnivory. *PLoS biology*, 5(2), 2007.

[130] Erin E Wilson and Elizabeth M Wolkovich. Scavenging: how carnivores and carrion structure communities. *Trends in Ecology & Evolution*, 26(3):129–135, 2011.

[131] Erika V Iyengar. Kleptoparasitic interactions throughout the animal kingdom and a re-evaluation, based on participant mobility, of the conditions promoting the evolution of kleptoparasitism. *Biological Journal of the Linnean Society*, 93(4):745–762, 2008.

[132] William K Steele and Philip AR Hockey. Factors influencing rate and success of intraspecific kleptoparasitism among kelp gulls (larus dominicanus). *The Auk*, 112(4):847–859, 1995.

[133] Chris Carbone, Georgina M Mace, S Craig Roberts, and David W Macdonald. Energetic constraints on the diet of terrestrial carnivores. *Nature*, 402(6759):286–288, 1999.

[134] Graeme D Ruxton and David C Houston. Obligate vertebrate scavengers must be large soaring fliers. *Journal of theoretical biology*, 228(3):431–436, 2004.

[135] Ester van der Meer, Mkhalalwa Moyo, Gregory SA Rasmussen, and Hervé Fritz. An empirical and experimental test of risk and costs of kleptoparasitism for african wild dogs (lycaon pictus) inside and outside a protected area. *Behavioral Ecology*, 22(5):985–992, 2011.

[136] DBA Thompson. The economics of kleptoparasitism: optimal foraging, host and prey selection by gulls. *Animal Behaviour*, 34(4):1189–1205, 1986.

[137] Martyn L Gorman, Michael G Mills, Jacobus P Raath, and John R Speakman. High hunting costs make african wild dogs vulnerable to kleptoparasitism by hyaenas. *Nature*, 391(6666):479–481, 1998.

[138] JS Hunter, SM Durant, and TM Caro. To flee or not to flee: predator avoidance by cheetahs at kills. *Behavioral Ecology and Sociobiology*, 61(7):1033–1042, 2007.

[139] Laura M Pereira, Norman Owen-Smith, and Marcos Moleón. Facultative predation and scavenging by mammalian carnivores: seasonal, regional and intra-guild comparisons. *Mammal Review*, 44(1):44–55, 2014.

[140] Daniel H Janzen. Why fruits rot, seeds mold, and meat spoils. *The American Naturalist*, 111(980):691–713, 1977.

[141] Deron E Burkepile, John D Parker, C Brock Woodson, Heath J Mills, Julia Kubanek, Patricia A Sobecky, and Mark E Hay. Chemically mediated competition between microbes and animals: microbes as consumers in food webs. *Ecology*, 87(11):2821–2831, 2006.

[142] John A Shivik. Are vultures birds, and do snakes have venom, because of macro-and microscavenger conflict? *Bioscience*, 56(10):819–823, 2006.

[143] Marcos Moleón, Nuria Selva, Maria Martina Quaggiotto, David M Bailey, Ainara Cortés-Avizanda, and Travis L DeVault. Carrion availability in space and time. In *Carrion ecology and management*, pages 23–44. Springer, 2019.

[144] Jerry A Payne. A summer carrion study of the baby pig sus scrofa linnaeus. *Ecology*, 46(5):592–602, 1965.

[145] H Jane Brockmann and CJ Barnard. Kleptoparasitism in birds. *Animal behaviour*, 27:487–514, 1979.

[146] K Nishimura. Kleptoparasitism and cannibalism. *MD Breed & J. Moore*, 2010.

[147] PJ Funston, MGL Mills, HC Biggs, and PRK Richardson. Hunting by male lions: ecological influences and socioecological implications. *Animal Behaviour*, 56(6):1333–1345, 1998.

[148] Warren B Ballard, Jackson S Whitman, and Craig L Gardner. Ecology of an exploited wolf population in south-central alaska. *Wildlife monographs*, pages 3–54, 1987.

[149] Travis L DeVault, Olin E Rhodes, Jr, and John A Shivik. Scavenging by vertebrates: behavioral, ecological, and evolutionary perspectives on an important energy transfer pathway in terrestrial ecosystems. *Oikos*, 102(2):225–234, 2003.

[150] Neil Hammerschlag, R Aidan Martin, and Chris Fallows. Effects of environmental conditions on predator–prey interactions between white sharks (carcharodon carcharias) and cape fur seals (arctocephalus pusillus pusillus) at seal island, south africa. *Environmental Biology of Fishes*, 76(2):341–350, 2006.

[151] Tobey H Curtis, John T Kelly, Karl L Menard, R Karl Laroche, Robert E Jones, and A Peter Klimley. Observations on the behavior of white sharks scavenging from a whale carcass at point reyes, california. *California Fish and Game*, 92(3):113, 2006.

[152] JP Croxall and PA Prince. Dead or alive, night or day: how do albatrosses catch squid? *Antarctic Science*, 6:155–155, 1994.

[153] Travis L DeVault and Olin E Rhodes. Identification of vertebrate scavengers of small mammal carcasses in a forested landscape. *Acta Theriologica*, 47(2):185–192, 2002.

[154] David C Houston. The adaptations of scavengers. *Serengeti: dynamics of an ecosystem*, pages 263–286, 1979.

[155] ML Wilton. Scavenging and its possible effects upon predation-a selective review of literature. *Alces*, 22:155–180, 1986.

[156] Hans Kruuk. The spotted hyena: a study of predation and social behavior. Technical report, 1972.

[157] Craig Packer, David Scheel, and Anne E Pusey. Why lions form groups: food is not enough. *The American Naturalist*, 136(1):1–19, 1990.

[158] Adam Kane, Kevin Healy, Thomas Guillerme, Graeme D Ruxton, and Andrew L Jackson. A recipe for scavenging in vertebrates–the natural history of a behaviour. *Ecography*, 40(2):324–334, 2017.

[159] Tom P Flower, Matthew F Child, and Amanda R Ridley. The ecological economics of kleptoparasitism: pay-offs from self-foraging versus kleptoparasitism. *Journal of Animal Ecology*, 82(1):245–255, 2013.

[160] RW Furness. Kleptoparasitism by great skuas (catharacta skua brünn.) and arctic skuas (stercorarius parasiticus l.) at a shetland seabird colony. *Animal Behaviour*, 26:1167–1177, 1978.

[161] Malte Andersson. Predation and kleptoparasitism by skuas in a shetland seabird colony. *Ibis*, 118(2):208–217, 1976.

[162] Karen R Cangialosi. Social spider defense against kleptoparasitism. *Behavioral Ecology and Sociobiology*, 27(1):49–54, 1990.

[163] C Carbone, JT Du Toit, and IJ Gordon. Feeding success in african wild dogs: does kleptoparasitism by spotted hyenas influence hunting group size? *Journal of animal Ecology*, pages 318–326, 1997.

[164] EV Iyengar. To steal or not to steal? that is the question. suspension feeding versus kleptoparasitism in a marine snail. In *AMERICAN ZOOLOGIST*, volume 40, pages 1073–1073. AMER SOC ZOOLOGISTS 1041 NEW HAMPSHIRE ST, LAWRENCE, KS 66044 USA, 2000.

[165] M Broom and GD Ruxton. Evolutionarily stable kleptoparasitism: consequences of different prey types. *Behavioral Ecology*, 14(1):23–33, 2003.

[166] Kevin A Blecha, Randall B Boone, and Mathew W Alldredge. Hunger mediates apex predator's risk avoidance response in wildland–urban interface. *Journal of Animal Ecology*, 87(3):609–622, 2018.

[167] Anja Petersen, Kristian T Nielsen, Christian B Christensen, and Søren Toft. The advantage of starving: success in cannibalistic encounters among wolf spiders. *Behavioral Ecology*, 21(5):1112–1117, 2010.

[168] Stan L Lindstedt and Mark S Boyce. Seasonality, fasting endurance, and body size in mammals. *The American Naturalist*, 125(6):873–878, 1985.

[169] Justin D Yeakel, Christopher P Kempes, and Sidney Redner. Dynamics of starvation and recovery predict extinction risk and both damuth's law and cope's rule. *Nature communications*, 9(1):1–10, 2018.

[170] Jenny Mattisson, Geir Rune Rauset, John Odden, Henrik Andrén, John DC Linnell, and Jens Persson. Predation or scavenging? prey body condition influences decision-making in a facultative predator, the wolverine. *Ecosphere*, 7(8):e01407, 2016.

[171] Annelies De Cuyper, Marcus Clauss, Chris Carbone, Daryl Codron, An Cools, Myriam Hesta, and Geert PJ Janssens. Predator size and prey size–gut capacity ratios determine kill frequency and carcass production in terrestrial carnivorous mammals. *Oikos*, 128(1):13–22, 2019.

[172] Adam Kane, Kevin Healy, Graeme D Ruxton, and Andrew L Jackson. Body size as a driver of scavenging in theropod dinosaurs. *The American Naturalist*, 187(6):706–716, 2016.

[173] Rudolf Philippe Rohr, Heike Scherer, Patrik Kehrli, Christian Mazza, and Louis-Félix Bersier. Modeling food webs: exploring unexplained structure using latent traits. *The American Naturalist*, 176(2):170–177, 2010.

[174] Justin D Yeakel, Mathias M Pires, Lars Rudolf, Nathaniel J Dominy, Paul L Koch, Paulo R Guimarães, and Thilo Gross. Collapse of an ecological network in ancient egypt. *Proceedings of the National Academy of Sciences*, 111(40):14472–14477, 2014.

[175] Mathias M Pires, Paul L Koch, Richard A Farina, Marcus AM de Aguiar, Sérgio F dos Reis, and Paulo R Guimarães Jr. Pleistocene megafaunal interaction networks became more vulnerable after human arrival. *Proceedings of the Royal Society B: Biological Sciences*, 282(1814):20151367, 2015.

[176] Colin W Clark, Marc Mangel, et al. *Dynamic state variable models in ecology: methods and applications*. Oxford University Press on Demand, 2000.

[177] John Damuth. Interspecific allometry of population density in mammals and other animals: the independence of body mass and population energy-use. *Biological Journal of the Linnean Society*, 31(3):193–246, 1987.

[178] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.

[179] Theodore M DeJong. A comparison of three diversity indices based on their components of richness and evenness. *Oikos*, pages 222–227, 1975.

[180] Evelyn C Pielou. The measurement of diversity in different types of biological collections. *Journal of theoretical biology*, 13:131–144, 1966.

[181] Laura R Prugh. Coyote prey selection and community stability during a decline in food supply. *Oikos*, 110(2):253–264, 2005.

[182] Benjamin L Allen. Skin and bone: observations of dingo scavenging during a chronic food shortage. *Australian Mammalogy*, 32(2):207–208, 2010.

[183] James D Roth. Variability in marine resources affects arctic fox population dynamics. *Journal of Animal Ecology*, 72(4):668–676, 2003.

[184] Stéphanie Périquet, Marion Valeix, Jaelle Claypole, Nolwenn Drouet-Hoguet, Julia Salnicki, Simeon Mudimba, Eloy Revilla, and Hervé Fritz. Spotted hyaenas switch their foraging strategy as a response to changes in intraguild interactions with lions. *Journal of Zoology*, 297(4):245–254, 2015.

[185] Laura R Prugh and Kelly J Sivy. Enemies with benefits: integrating positive and negative interactions among terrestrial carnivores. *Ecology Letters*, 23(5):902–918, 2020.

[186] Mark D Bertness and Ragan Callaway. Positive interactions in communities. *Trends in ecology & evolution*, 9(5):191–193, 1994.

[187] Oliver P Höner, Bettina Wachter, Marion L East, and Heribert Hofer. The response of spotted hyaenas to long-term changes in prey populations: functional response and interspecific kleptoparasitism. *Journal of Animal Ecology*, 71(2):236–246, 2002.

[188] Susan M Cooper, Kay E Holekamp, and Laura Smale. A seasonal feast: long-term analysis of feeding behaviour in the spotted hyaena (crocuta crocuta). *African Journal of Ecology*, 37(2):149–160, 1999.

[189] SM Cooper. Optimal hunting group size: the need for lions to defend their kills against loss to spotted hyaenas. *African Journal of Ecology*, 29(2):130–136, 1991.

[190] Jürgen Jacobs. Quantitative measurement of food selection. *Oecologia*, 14(4):413–417, 1974.

[191] Matt W Hayward, John O'Brien, Markus Hofmeyr, and Graham IH Kerley. Prey preferences of the african wild dog lycaon pictus (canidae: Carnivora): ecological requirements for conservation. *Journal of Mammalogy*, 87(6):1122–1131, 2006.

[192] MW Hayward, M Hofmeyr, J O'brien, and Graham IH Kerley. Prey preferences of the cheetah (acinonyx jubatus)(felidae: Carnivora): morphological limitations or the need to capture rapidly consumable prey before kleptoparasites arrive? *Journal of Zoology*, 270(4):615–627, 2006.

[193] MW Hayward, Philippe Henschel, John O'Brien, Markus Hofmeyr, Guy Balme, and Graham IH Kerley. Prey preferences of the leopard (panthera pardus). *Journal of Zoology*, 270(2):298–313, 2006.

[194] MW Hayward. Prey preferences of the spotted hyaena (crocuta crocuta) and degree of dietary overlap with the lion (panthera leo). *Journal of Zoology*, 270(4):606–614, 2006.

[195] Matt W Hayward and Graham IH Kerley. Prey preferences of the lion (panthera leo). *Journal of zoology*, 267(3):309–322, 2005.

[196] MW Hayward, W Jedrzejewski, and B Jedrzejewska. Prey preferences of the tiger p anthera tigris. *Journal of Zoology*, 286(3):221–231, 2012.

[197] Ch Bissett and RTF Bernard. Habitat selection and feeding ecology of the cheetah (acinonyx jubatus) in thicket vegetation: is the cheetah a savanna specialist? *Journal of Zoology*, 271(3):310–317, 2007.

[198] Guy Balme, Luke Hunter, and ROB Slotow. Feeding habitat selection by hunting leopards panthera pardus in a woodland savanna: prey catchability versus abundance. *Animal Behaviour*, 74(3):589–598, 2007.

[199] Susana Rostro-García, Jan F Kamler, and Luke TB Hunter. To kill, stay or flee: the effects of lions and landscape factors on habitat and kill site selection of cheetahs in south africa. *PloS one*, 10(2):e0117743, 2015.

[200] John H Fanshawe and Clare D Fitzgibbon. Factors influencing the hunting success of an african wild dog pack. *Animal behaviour*, 45(3):479–490, 1993.

[201] Kay E Holekamp, Laura Smale, R Berg, and Susan M Cooper. Hunting rates and hunting success in the spotted hyena (crocuta crocuta). *Journal of Zoology*, 242(1):1–15, 1997.

[202] Stan L Lindstedt and PJ Schaeffer. Use of allometry in predicting anatomical and physiological parameters of mammals. *Laboratory animals*, 36(1):1–19, 2002.

[203] Henry D Prange, John F Anderson, and Hermann Rahn. Scaling of skeletal mass to body mass in birds and mammals. *The American Naturalist*, 113(1):103–122, 1979.

[204] Geoffrey B West, James H Brown, and Brian J Enquist. A general model for ontogenetic growth. *Nature*, 413(6856):628–631, 2001.

[205] Chen Hou, Wenyun Zuo, Melanie E Moses, William H Woodruff, James H Brown, and Geoffrey B West. Energy uptake and allocation during ontogeny. *science*, 322(5902):736–739, 2008.

[206] Ewald R Weibel, Leonardo D Bacigalupe, Beat Schmitt, and Hans Hoppeler. Allometric scaling of maximal metabolic rate in mammals: muscle aerobic capacity as determinant factor. *Respiratory physiology & neurobiology*, 140(2):115–132, 2004.

[207] Matteo Rizzuto, Chris Carbone, and Samraat Pawar. Foraging constraints reverse the scaling of activity time in carnivores. *Nature ecology & evolution*, 2(2):247–253, 2018.

[208] Samraat Pawar, Anthony I Dell, and Van M Savage. Dimensionality of consumer search space drives trophic interaction strengths. *Nature*, 486(7404):485–489, 2012.