

Evolving Perceptual Categories

Cailin O'Connor

Department of Logic and Philosophy of Science

University of California, Irvine

`cailino@uci.edu`

1 Introduction

Hume's scepticism was based, in part, on what he believed to be an inability of human observers to ever go beyond their perceptual experiences. How can an observer determine the accuracy of perception when her only access to the external world is mediated by perception itself? A related question may be asked about not the *qualities* of perception, but rather the perceptual *categories* of biological actors. How can one reach beyond perception to determine whether and to what degree perceptual categories—sweet, green, cool—track properties of the world around them?

Jäger (2007) introduced a set of games—‘sim-max games’—that can help address this problem. Sim-max games are modified versions of the signaling game, introduced by Lewis (1969), that assume actors are transferring information about states of the world that bear underlying similarity relations to one another. What this means is that, unlike in a traditional Lewis signaling game, there is a natural sense in which actors in sim-max games might want to categorize groups of states together in order to transfer meaning about them.¹

This paper will use sim-max games to model perceptual categorization, with the goal of answering a question related to the problem mentioned above: to what degree should we expect the perceptual categories of biological actors to track properties of the world around them? I will argue that an analysis of these games suggests that the relationship between real-world structure and evolved perceptual categories is mediated by successful action in the sense that organisms evolve to categorize together states of nature for which similar actions lead to similar results.

¹Jäger is not the first to introduce this type of structure to the state space of a signaling game. Spence (1973), in his famous paper on job market signaling, does so too. Spence's model does not assume complete common interest between the actors, and for this reason Jäger's models are used here.

This conclusion indicates that both strongly realist and strongly anti-realist views about perceptual categories are too simple.

Before beginning, I would like to make a methodological remark. One might worry that the central observation of this paper—that perceptual categories can be expected to group real-world items for which the same actions are effective—could be made without appeal to game theory. One can give an intuitive argument for this observation; evolution responds to payoff and thus perceptual categories will evolve to track payoff rather than natural structure. And, in fact, both vision researchers and philosophers of color have made arguments along these lines.^{2,3} The game theoretic framework presented here, though, does two things that this intuitive argument does not. First, sim-max games provide mathematical justification for what might otherwise seem to be a hand-wavy conclusion. Second, the framework brings conceptual clarity to a topic that is complex and many faceted. While the thesis of the present paper may seem obvious, or even trivial, from a strategic point of view, it is certainly *not* universally accepted by philosophers or scientists studying perception.^{4,5}

The paper will proceed as follows. In section 2, I describe sim-max games. In section 3, I outline optimal strategies for these games and discuss their evolutionary properties. In section 4, I describe how these games can be used as a model of the evolution of perceptual categories and then proceed to make the principal arguments of the paper.

2 Signaling Games and Sim-Max Games

The standard signaling game, as outlined by Lewis (1969), is often taken as a model of information transfer between agents. The game has two players—a sender and a

²In their “wholly empirical theory of perception” Dale Purves and other vision researchers are interested in explaining how the visual system solves the inverse optics problem—that a single visual array can be caused by an infinite number of three dimensional stimuli. On their theory, this occurs through evolutionary and learning processes that lead to successful (rather than correct) interpretations of sensory input. Purves et al. (2011) argue that, “...the basis for what we see is not the physical qualities of object or actual conditions in the world but operationally determined perceptions that promote behaviors that worked in the past and are thus likely to work in response to current retinal stimuli” (1).

³Thompson (1995) argues that the function of color vision is to divide surfaces into a small set of color categories that are useful as a guide to behavior.

⁴See, for example, Mark et al. (2010) and Hoffman (2009) for insight into standard views in perceptual science. They argue that most perceptual researchers hold the ‘conventional view’ of perceptual categories—that they veridically track real-world structure. See Marr (1982) as an example of the type of position Hoffman et al. have in mind.

⁵See work by David Hilbert and Mohen Matthen for example of prominent philosophical positions that do not accept this conclusion.

receiver—and three stages. In stage one, exogenous forces, or ‘nature’, determine the state of the world. In stage two, the sender observes this state and sends a signal to the receiver contingent on it. In stage three, the receiver observes the sent signal and chooses an action. If the action is appropriate for the state of the world, both players receive a payoff. If the action is inappropriate, neither player receives a payoff. The goal of both players is to coordinate the action taken by the receiver with the state observed by the sender. Neither player cares how this coordination is achieved, i.e., what specific signal is used to denote any particular state, as long as coordination occurs.

Sim-max games build on this model by adding similarity structure to the state space of the signaling game. In their basic form, these games model situations where organisms would like to transfer information about properties that vary finely, or even continuously. In these cases, because there are many relevant states of the world that bear similarity relations to one another, categorization is potentially useful for transferring information. Properties that fit this description are things like distance, size, degree, concentration, color, time duration, temperature, etc.

In sim-max games, states of the world are modeled as points in a metric space where distance represents similarity.⁶ The greater the distance between two states, the less similar they are and vice versa. In these games, as in the Lewis signaling game, ‘nature’ selects a state of the world. The sender observes this state and sends a signal. The receiver observes this signal and chooses an act. If the act is perfectly appropriate for the state of the world, the actors receive identical, perfect payoffs. However, the key alteration to the game is that if the act is not perfectly appropriate, the actors still receive a payoff based on how nearly appropriate the act is. In other words, the game incorporates payoffs that vary as a function of the distance between the state of the world and the act taken. Specifically, payoff decreases as distance between the state of the world and the act taken increases. This aspect of the model mimics many real-world situations where states bear similarity relations to one another. In these cases, because states are similar, the same action will often be appropriate (or nearly appropriate) for a number of them. If we are signaling about the ripeness of an apple, for example, I might group a number of apples with different intensities of redness under the signal ‘ripe’. If you bite into any one of these, you will receive a payoff because the apples are ripe or close to it. This will be the case even if some of the redder apples are slightly more ripe than your ideal apple, and some of the greener ones are slightly less so.

Before continuing, it will be useful to say a word about similarity, given the importance of this concept to the sim-max model. Similarity is a complex and

⁶For simplicity sake, this paper will always assume that this space is a subset of \mathbb{R}^n endowed with a Euclidean metric.

much debated notion. In the most clear cut cases, the claim that real-world states can be more or less similar is perhaps not contentious. Consider a situation in which the relevant states of the world correspond to the length of a bridge. It is not strange to say that a 200 foot bridge is more like a 201 foot bridge than a 4 foot bridge holding all other features constant. In cases like this, one is considering similarity with respect to some feature of objects.⁷ Moving away from these cases, one quickly runs into confusing situations. Is a banana objectively more like a skyscraper or a wink? For the purposes of this paper, problems regarding similarity will have to be set aside. It will be assumed that there *are* ways in which real-world states are similar and dissimilar. The lessons drawn about relationships between perceptual categories and real-world similarity, though, should be taken to apply more clearly in less contentious cases of similarity.⁸

3 Optimality, Equilibria, and Evolution

Having outlined sim-max games, it will now be useful to discuss these games in more detail in order to understand how they shed light on perceptual categorization. In particular, I will discuss optimality and some results from evolutionary game theory that will help elaborate what sorts of strategies should be expected to evolve when actors are playing sim-max games.

3.1 Strategies and Optimality in Sim-Max Games

Optimal strategies in sim-max games are those strategies where similar states are categorized under the same signal, and where the actor then takes an action in response to that signal that is at least partially appropriate for all the states attached to it. In other words, the receiver is taking acts that are as appropriate as possible as much of the time as possible. Such a strategy will always garner the highest payoffs for the actors due to the assumption that payoff decreases with distance. The question, then, is which strategies allow the actors to most perfectly coordinate action?

To understand these strategies, it will first be useful to describe what is called a *Voronoi tessellation*. A Voronoi tessellation is a division of a space around what are sometimes called ‘generators’ or ‘seeds’. These generators consist of locations

⁷Goodman (1972) argues that similarity only makes sense with respect to some feature.

⁸It should be noted that the model presented in this section is far from the first formal model of categorization. There is a very large formal literature in psychology and cognitive science related to categorization. This work mainly attempts to model experimental results on human linguistic categorization. For an overview of this literature see Kruschke (2008) or Murphy (2002). There is a smaller categorization literature in economics that focuses more on optimal behavior. See Mohlin (2013) as an example.




Figure 1: Examples of Voronoi tessellations. Diagrams (a) and (b) show tessellations of a line. Diagrams (c) and (d) show tessellations of a two dimensional space.

in the space. The tessellation is a division of the surrounding space into cells where every point is assigned to a cell based on which generator it is closest to. Figure 1 shows four Voronoi tessellations. The first two are of a straight line. Each point pictured in diagrams (a) and (b) is a generator, and each cell of the tessellation is represented by two vertical dashes on either side of these generators. Note that in each case the vertical dashes are exactly equidistant from the generators on either side of them. The diagrams (c) and (d) show Voronoi tessellations of a two dimensional space. Each generator is represented as a point, and each cell contains only those points in the space closest to the generator at the center of that cell.

Suppose that the space in diagram (d) represents the state space of a sim-max game. Furthermore, suppose that each point represents the ideal state for an act that a receiver is taking in response to some signal. The cells of the Voronoi tessellation represent the optimal response by the sender to the receiver. For any particular state of the world, the sender assigns that state a signal based on which cell it is a part of. Given which acts the receiver is taking for each signal, this sender strategy ensures that each state will garner the highest payoff possible, because it will be matched with the closest possible act. Given this type of optimal sender strategy, receivers then optimize by spacing out signal interpretations so that the Voronoi cells associated with each one will be nearly the same size.⁹ Diagrams (b)

⁹This is only true when states are equiprobable and the payoff function is the same over all states of the world. In reality, this will never be the case. What this observation should be taken to show is roughly that signals should be allocated to cover regions of states that are of equal importance from a payoff perspective.

and (c) represent Voronoi strategies with this sort of optimal character.^{10,11,12}

3.2 Evolution and Emergence

What has not yet been addressed, but will be important to subsequent discussion, is what the optimality of these strategies means from an evolutionary perspective. As it turns out, there are existing results from evolutionary game theory that can help inform the problem at hand.

The first thing to note is that the optimal strategies of sim-max games are always what are called payoff dominant Nash equilibria of the game. A Nash equilibrium of a game is a set of strategies from which no player can unilaterally deviate and improve her payoff. Payoff dominant Nash equilibria are those where the actors achieve the highest possible payoffs. These equilibria are significant from an evolutionary point of view. The replicator dynamics—the most widely studied model of evolutionary change in evolutionary game theory—work by successively increasing the frequency of strategies in a population of players that get high payoffs and decreasing the frequency of strategies that get low payoffs. Thus, strategies that are more successful for both actors are often those that evolve. In standard Lewis signaling games it has been shown that the payoff dominant equilibria of the game, called ‘signaling systems’, will evolve under the replicator dynamics in many cases.¹³ It is not the case, however, that populations investigated using the replicator dynamics *always* evolve to states where the actors play payoff dominant strategies. In fact, for Lewis signaling games, it has also been shown that (except in the special case of the 2 state, 2 signal, 2 act game where states are equally likely) populations can evolve to play less efficient strategies.¹⁴ Given these observations, it cannot be assumed that players in models using sim-max games will always evolve the type of optimal strategies that have been described here. What *should* they be expected to evolve?

Jäger (2007), in an analysis of sim-max games and their behavior under the replicator dynamics, has shown that the asymptotically stable rest points of the game will be those where, “the sender strategy is consistent with the Voronoi tessel-

¹⁰Only pure strategies of these games are considered here as Lipman (2009) has shown that in common interest signaling scenarios mixed strategies never outperform pure strategies and in all but a few ‘knife’s edge’ cases do strictly worse.

¹¹This is obviously a very informal description of these optimal strategies. For a more detailed treatment, see Jäger et al. (2011). Jäger (2007) discusses optimal sender strategies of discrete sim-max games. O’Connor (2013) discusses optimal strategies in discrete games with a one dimensional state space at greater length.

¹²Mohlin (2013) has found similar optimality results in a model of categorization for the purposes of prediction. In Mohlin’s model, unlike here, optimal category size emerges endogenously.

¹³See, for example, Skyrms (1996), Skyrms (2000), and Huttegger et al. (2010).

¹⁴See Huttegger et al. (2010), Huttegger (2007), and Pawlowitsch (2008).

lation that is induced by the image of the receiver strategy” (562). In other words, under the replicator dynamics, categories—as represented by signaler strategies—that Voronoi tessellate the space will be the asymptotically stable rest points. An asymptotically stable rest point under the replicator dynamics is a state where a population, if perturbed away from this state, will return. The stability of these rest points makes them significant from an evolutionary point of view. Furthermore, these categories will be ideally suited to respond to the evolved receiver strategy. This holds for the symmetrized version of the game where each actor has both a sender and a receiver strategy.

Jäger says nothing about what receiver strategy should be expected to evolve in sim-max games, besides restricting them to pure strategies. In fact, though, the strategy set that he has delineated as those that might evolve can be greatly restricted. As Jäger has argued, the strategies that will evolve are those where the sender strategy is a Voronoi tessellation of the space in response to the receiver strategy. However, there are some cases where this is the case, but where the receiver strategy is not the best response to the sender strategy. In other words, the receiver can unilaterally deviate and improve payoff, which means that the strategies played are not part of a Nash equilibrium. Non-Nash strategies are never stable under the replicator dynamics, and so will not evolve. The restriction of Jäger’s strategies to Nash equilibria will include all the optimal strategies. It will also include some strategies that are non-optimal, but, intuitively, are still strategies where the sender divides the state space nearly evenly into categories and the receiver responds appropriately.^{15,16}

Jäger takes his results to show that Voronoi strategies will *always* be those that evolve for sim-max games under the replicator dynamics, but recent work by Elliott Wagner complicates these results. Wagner (personal correspondence) has shown that certain non-Voronoi strategies will be Lyapunov stable for the symmetrized version of some sim-max games. Lyapunov stable rest point is one for which a population near that rest point will stay near it. All asymptotically stable rest points are Lyapunov stable, but the reverse is not true. In this sense, asymptotic stability is a stronger type of stability. In other words, in certain cases, non-Voronoi strategies can evolve. Wagner’s results do indicate, however, that in

¹⁵It should be noted that a subset of strategies called ‘babbling’ strategies will both be Nash equilibria and Voronoi tessellations in response to receiver strategies, but will not have this character of near-optimality. In babbling strategies, the sender always sends the same signal no matter the state, and the receiver always takes the same act no matter the signal. These strategies are not asymptotically stable in sim-max games, however. In particular, they will not be part of what is called an evolutionarily stable set (ESSet). Cressman (2003) has shown that for games of the type considered here, the asymptotically stable rest points of the game are all and only those strategies that are part of ESSets.

¹⁶Jäger et al. (2011) has similar results for games where the state space is continuous.

cases where states of the world are numerous, non-Voronoi strategies are less likely to evolve, so it may be that in models that closely correspond to real world cases (where states of the world tend to be many) Jäger is largely correct.

4 Applying Sim-Max Games to Perceptual Categorization

Signaling games are most often taken as models of information transfer *between* organisms. In fact, though, signaling games can also be used to understand *within* organism processes. In particular, signaling games can be interpreted as representing perceptual signaling. Under this interpretation, the state of the world is a type of real-world item that causes perceptual input while the act taken is the reaction of the organism to this perceptual input. It is difficult to say exactly what, in this scenario, the signal is. However, if one allows for a broad interpretation, the signal can be taken to be a mediating mental process that starts with perceptual input and ends with behavior. In particular, one can take the perceptual experience of an organism to be an important part of this signal.¹⁷

With this interpretation, one can use sim-max games as a model of perceptual categorization. Upon doing so, the first thing to note is that the evolved categories in sim-max games—Wagner’s results notwithstanding—are convex. A convex set in a metric space is one where for any two points in the set, every point on a straight line segment between these two points is in the set as well. The cells of Voronoi tessellations are necessarily convex, and so the categories that evolve in the long run in sim-max games are convex as well. Convexity is interesting here because it means that actors evolving perceptual categories should be expected to group together states of the world in ways that reflect what are arguably natural properties of these states. Like things will be grouped together, sufficiently unlike things will not. In other words, perceptual categories should be expected to track real-world structure in that similar real-world items will belong to the same category while sufficiently different real-world items will not.

However, a closer look at just how similarity is built into the structure of sim-max games indicates that, in fact, this conclusion is not warranted. In describing sim-max games, I made an appeal to the geometrical structure of the state spaces of the games. Throughout the paper, to this point, distance in this space has been described as representing similarity between different states of the world. This

¹⁷In a somewhat related interpretation, Jäger (2007) uses sim-max games to understand ‘conceptual spaces’. He interprets the states of the game as corresponding to perceptual states, rather than real world states, and argues that these games provide evidence that concepts map to convex regions of perceptual space, as argued by Gardenfors (2000). His interpretation assumes that similar perceptual states will get similar payoffs in response to action. While he does not argue for this assumption, the results provided here indicate that it may be warranted.

is a natural way to describe these games and to give an intuitive sense of how they work and why they behave as they do. But, in fact, this intuitive way of understanding sim-max games is somewhat misleading. Similarity relations in the state space of sim-max games are built in through *payoff structure* alone. In other words, there is no necessary sense in which two states are like or unlike each other except in that the same acts will receive payoff for both states, or will not. The game can be described perfectly, without appealing to geometry, as a set of states, a set of acts, and a payoff defined for each pair of these. Another way to put this is as follows: from a modeling perspective, two states, even if they are objectively wildly different, should be treated as occupying the same spot in state space in a sim-max game if the same acts obtain the same payoffs in both states.

What this means is that results from sim-max games only indicate the following: organisms will evolve to categorize together states of the world for which the same act is appropriate, whether or not these states are somehow similar or dissimilar. Organisms, then, should be expected to evolve perceptual categories not directly on the basis of how real-world properties relate, but on the basis of how organism action either contributes or does not contribute to the organism's success when these properties are present.¹⁸ Another way of saying this might be that perceptual categories should track payoff similarity rather than real-world similarity.

The observations above, at first glance, may seem to strike a significant blow to those committed to the idea that perceptual categories accurately track real-world properties. Sim-max games seem to indicate that there is no particular reason to think this is the case. Mark et al. (2010) have argued something similar on the basis of a different set of evolutionary game theoretic models. They have shown that in competitive choice situations, perceptual categorization strategies that do not accurately mimic real-world structure can outperform those that do. Along these lines, Hoffman claims that perceptual categories should not be assumed to accurately reflect real-world properties because, as he says, "Fitness, not accuracy, is the *objective function* optimized by evolution" (5, emphasis his).

Such considerations might lead one to a strong anti-realist position with regards to perceptual categories—that there is no reason to think they track natural properties—and this is indeed what Hoffman defends.¹⁹ But sim-max games actually support only a more modest claim. According to the sim-max game framework, perceptual categorizations are determined through payoff constraints rather

¹⁸Barrett (2007) employs somewhat similar games and draws a somewhat related conclusion: that the dispositions of the actors in his games importantly determine how they will partition states of nature for the purposes of signaling. From this he concludes that language cannot be assumed to track natural kinds.

¹⁹Hoffman actually takes an even stronger anti-realist position which he calls the 'Interface Theory of Perception'. According this theory, perceptual categories specifically *do not* track real-world structure. See Mark et al. (2010) and Hoffman (2009).

than directly by real-world properties. A strong anti-realist position requires an additional assumption, however—that there is little relation between real-world structure and payoff structure. And this assumption is unwarranted.

It seems likely, for example, that when organisms encounter two highly similar real-world situations, the same actions will be appropriate to both. Take two scenarios which are identical with the exception that in the first a leopard is 20 feet away from a vervet and in the second the leopard is 21 feet away. Because the physical realities of the two scenarios are similar, it will be appropriate for the vervet to take the same action in both cases. In many cases, similar real-world states should, in sim-max games, be mapped to the same area of state space, because they are appropriately paired with the same sets of actions. If so, it should then be expected that perceptual categories track real-world similarity properly.

It is interesting to note, however, that in many situations, the same cannot be said about the differentiation of unlike states. It is often the case that similar actions will be appropriate for states of the world that are objectively different. If an escape pattern for vervets works both to avoid, say, leopards *and* human hunters, there is no reason not to use the same category to denote these two quite different groups of states. This argument about categorization is well embodied by the example of human color perception. It is notoriously difficult to identify some objective property of the world that color experience perfectly corresponds to. The most promising property is the surface spectral reflectance profile, or SSR, of a surface. This profile is a measurement of the probabilities with which wavelengths of light will reflect from a surface. While nearly identical SSRs will stimulate similar color experiences in observers, sometimes quite different SSRs will also stimulate similar color experiences in observers. In other words, color categorization respects similarity, but not necessarily difference.

5 Conclusion

To conclude, I would like to reemphasize the two main insights that the strategic framework presented here can provide with regard to perceptual categories. 1) Perceptual categories should be expected to evolve such that states that can be responded to with the same actions are grouped together perceptually and states that cannot are not. 2) The fact that evolutionary processes are guided by payoff does not necessarily mean that perceptual categories will fail to track real-world structure. Neither, though, do these results support what Hoffman (2009) calls the ‘conventional view’ of perception, that, “The primary goal of perceptual categorization is to recover, or estimate, the objective statistical structure of the physical world” (3). Instead, the sim-max game framework provides a better understanding

of how perceptual categories should be expected to relate to real-world structure based on how real-world structure influences organism payoff. This relationship means that perceptual categories should not necessarily be expected to either track or not track real-world structure, but rather to bear more complicated relationships to it.

References

- Barrett, Jeffrey A. (2007). “Dynamic Partitioning and the Conventionality of Kinds.” *Philosophy of Science*, 74(527-546).
- Cressman, R. (2003). *Evolutionary dynamics and extensive form games*. Cambridge, MA: MIT Press.
- Gärdenfors, Peter (2000). *Conceptual Spaces: on the geometry of thought*. Cambridge, MA: MIT Press.
- Goodman, N. (1972). “Seven Strictures on Similarity.” *Problems and Projects*. Ed. N. Goodman. New York: Bobs Merrill.
- Hoffman, Donald (2009). “The interface theory of perception: Natural selection drives true perception to swift extinction.” *Object categorization: Computer and human vision perspectives*. Ed. S. Dickinson, M. Tarr, A. Leonardis, and B. Schiele. Cambridge, UK: Cambridge University Press, 148–165.
- Huttegger, Simon M. (2007). “Evolution and the explanation of meaning.” *Philosophy of Science*, 74, 1–27.
- Huttegger, Simon M., Brian Skyrms, Rory Smead, and Kevin J. S. Zollman (2010). “Evolutionary dynamics of Lewis signaling games: signaling systems vs. partial pooling.” *Synthese*, 172, 177–191.
- Jäger, Gerhard (2007). “The evolution of convex categories.” *Linguistics and Philosophy*, 30, 551–564.
- Jäger, Gerhard, Lars P. Metzger, and Frank Riedel (2011). “Voronoi languages: Equilibria in cheap-talk games with high-dimensional types and few signals.” *Games and Economic Behavior*, November, 2011, 73(2), 517–537.
- Kruschke, John K. (2008). “Models of Categorization.” *The Cambridge Handbook of Computational Psychology*. Ed. R. Sun. Cambridge, UK: Cambridge University Press, 267–301.
- Lewis, David K. (1969). *Convention*. Cambridge, MA: Harvard University Press.
- Lipman, Barton (2009). “Why is Language Vague?.” Working paper, Boston University, Dept. of Economics, version of November 2009.
- Mark, Justin T., Brian B. Marion, and Donald Hoffman (2010). “Natural selection and veridical perceptions.” *Journal of Theoretical Biology*, 266(504-515).

- Marr, David (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: W. H. Freeman
- Mohlin, Erik (2013). “Optimal Cateogrization.”
- Murphy, G.L. (2002). *The Big Book of Concepts*. Cambridge, MA: MIT Press.
- O’Connor, Cailin (2013). “The Evolution of Vagueness.” *Erkenntnis*, 10.1007/s10670-013-9463-2.
- Pawlowitsch, Christina (2008). “Why evolution does not always lead to an optimal signaling system.” *Games and Economic Behavior*, 63, 203–226.
- Purves, Dale., William T. Wojtach, and R. Beau Lotto (2011). “Understanding vision in wholly empirical terms.” *Proceedings of the National Academy of Science*, 108(3)(15588-15595)..
- Skyrms, B. (1996). *Evolution of the Social Contract*. Cambridge: Cambridge University Press.
- Skyrms, B. (2000). “Stability and explanatory significance of some simple evolutionary models.” *Philosophy of Science*, 67(94-113).
- Spence, Michael (1973). “Job Market Signaling.” *The Quarterly Journal of Economics*, 87(3), 355–374.
- Thompson, Evan (1995). “Color Vision, Evolution, and Perceptual Content.” *Synthese*, 104, 1–32.