

Chapter

20

## Two “Inhibitions of Return” Bias Orienting Differently

Raymond M. Klein and Ralph S. Redden

### Discovery of an Inhibitory Aftereffect Following an Uninformative Peripheral Cue

An inhibitory aftereffect of orienting was described thoughtfully and explored thoroughly by Posner and Cohen (1984). Cited more than 3,000 times,<sup>1</sup> the “components of visual orienting” spawned a large, exciting, and multidisciplinary effort to measure the aftereffect’s spatial and temporal properties, determine its causes and effects, understand its functional significance, uncover its neural underpinnings, and explore the effects of psychoactive agents, neurological conditions, and various experimental manipulations on it.

In a prototypical experiment the participant fixates a central stimulus amid two or more equi-eccentric locations, marked by unfilled boxes. Then, the following sequence of events unfolds: An uninformative stimulus, often called a cue, is presented by brightening one of the peripheral boxes; at various times after the cue’s appearance, a target is presented randomly in one of the peripheral boxes. The participant is instructed to make a simple manual response as soon as the target is detected while remaining fixated on the central stimulus. Eye position is monitored by some means, and trials on which fixation is not maintained are excluded from analysis. In the original experiments one of two means was used to ensure that attention – assumed to have been captured by the spatially uninformative cue – would be returned to fixation: either targets were more likely at fixation than at peripheral locations or the central location was cued (i.e., there was a “cue-back”) prior to the appearance of the target. In the terminology first put forward by Posner (1980), these methods encourage withdrawal of attention from the cued location. When targets are probable at fixation, we assume that participants voluntarily disengage attention from the cued location and return it to fixation, a mode of control we refer to as endogenous. When the cue-back to fixation is used, we assume that attention is returned to fixation reflexively, a mode of control we refer to as exogenous.

The pattern of results that is typically obtained is biphasic.<sup>2</sup> When the interval between the cue’s and target’s onsets (the cue-target onset asynchrony, or CTOA) is brief, reaction time (RT) to targets at the cued location is faster than to targets at the uncued location. This facilitation of cued targets is presumed to be due to the capture of attention by the cue. When the CTOA is longer, assuming attention is removed from the cued location, RT to

<sup>1</sup> Google Scholar, April 2016.

<sup>2</sup> But not always. For example, failure to disengage attention from the cue may allow attention-based facilitation to overshadow the inhibitory aftereffect. Or, in the presence of an especially rapid disengagement, the early facilitation might not be seen. And, in some cases, early sensory adaptation (what some call “onset detection cost”) may overshadow attention-based facilitation.

targets at the cued location is slower than to targets at the uncued location, which reflects a presumed inhibitory aftereffect generated by the cue.

Posner and Cohen's seminal paper described many experiments that allowed quite a few inferences to be made about this inhibitory aftereffect, most of which have been substantiated in the 30 years of follow-up research: it could be generated by increasing or decreasing the luminance of a peripheral box; it is not due to guessing, attentional momentum, masking, or inhibiting a response to the cue; it is coded, not in retinotopic but in environmental coordinates; it is independent of and can coexist with facilitation; it is relatively long lasting; and it is not generated following a spatially informative, centrally presented arrow (endogenous orienting).

Posner and Cohen (1984) identified the cause of inhibition as any peripheral stimulation rather than the exogenous orienting normally generated by peripheral cues. We believe that this may be the only incorrect inference that they drew.<sup>3</sup> They identified the effect of inhibition as one that reduces "the efficiency of target detection." And they proposed that the function of the inhibition was to "maximize sampling of the visual environment."

## Inhibition of Return

In the next year, Posner, Rafal, Choate, and Vaughan (1985) explored an inhibitory aftereffect, which they called "inhibition of return" (IOR). In Experiment 1, they tested 4 patient groups using the model task pioneered by Posner and Cohen. Patients with progressive supranuclear palsy, which includes damage to the superior colliculus, a brain region that plays a critical role in exogenously generated saccades, failed to show an inhibitory aftereffect in the direction in which their saccadic behavior was compromised (vertically). In two subsequent experiments, Posner and colleagues made important discoveries about the causes and effects of IOR, discoveries that in retrospect diverged strikingly from what was described by Posner and Cohen.

In Experiment 2, IOR was generated using a version of the model task described earlier, but its effect was measured using a temporal order judgment task with two different response methods. On every trial, a target appeared at each of two locations with a range of short temporal intervals between their onsets. In one condition, participants made a temporal order judgment with a button press signaling on which side the first target appeared – the predicted inhibition did not affect temporal order judgment (a finding later replicated by Gibson & Egeth, 1994). In the other condition, participants were asked instead to make an eye movement in the most comfortable direction after the targets had appeared; here, there was a small bias against making eye movements in the cued direction: "The bias in eye movements is less than 10% at critical intervals, suggesting that it is not the major

---

<sup>3</sup> Two findings led to this inference: (a) Following a double cue (one on each side of fixation) that would not generate orienting, they found a delay (compared to uncued target targets from single-cue trials), and (b) there was no inhibition following a shift of attention generated endogenously. The first finding seems to show that you can get the inhibition without a shift of attention, while the second shows that the inhibition does not necessarily follow a shift of attention. We now know that inhibition is not generated by stimulation, because a balanced array of cues around fixation does not generate inhibition at cued relative to uncued locations (Klein, Christie, & Morris, 2005). We also know, or at least we believe, that exogenous and endogenous attention shifts are different in many ways (e.g., see Klein, 2009 for a review).

determiner of where the eyes will go, but rather a bias built on top of other factors which affect the eye movement systems” (Posner et al., 1985, p. 223).

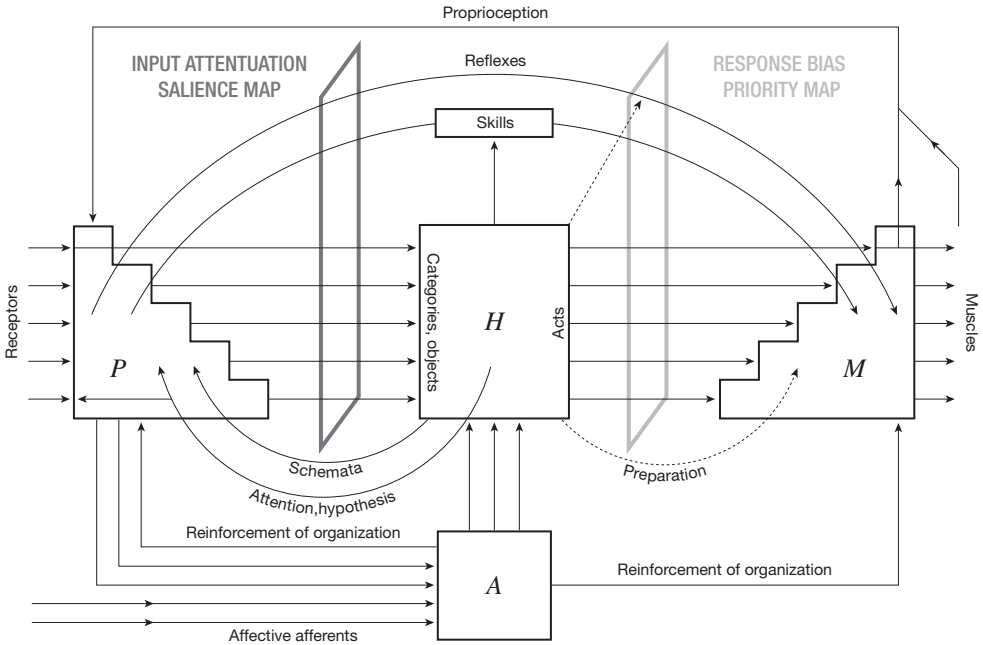
In Experiment 3, Posner and colleagues assiduously avoided the asymmetric stimulation of the two peripheral locations, which can generate orienting and oculomotor activity reflexively. Two small digits (too small to resolve in peripheral vision) were presented simultaneously, one in each peripheral box. Thus, unlike the more traditional single-cue experiments, the luminance energy at both locations was equivalent. Shortly afterward, an arrow at fixation called for the participants to move their eyes to the indicated box and report the digit. A cue back to fixation called for gaze and attention to return to the central fixation stimulus. One second later, a peripheral target calling for a manual detection response was presented at the previously fixated or un-fixated peripheral location. Manual detection response times were slower for targets presented at the previously fixated location. This inhibitory aftereffect was inferred to have been caused by the eye movement, because there had been no asymmetric stimulation of the periphery. Reinforcing the novelty-seeking attribution proposed by Posner and Cohen, Posner and colleagues hypothesized that IOR biases attention during the search of complex scenes, a proposal that has since been robustly confirmed (for a review, see Wang & Klein, 2010).

## Two Different Inhibitory Biases

Over the past 15 years, findings have emerged to support the implication, which goes back to these two seminal studies, that there are two different forms of IOR. We believe that both forms are caused by activity in or reaching the oculomotor system; both forms bias orienting, and hence, both forms can perform the novelty-seeking function attributed to the inhibition(s) in Posner’s seminal papers. Despite these shared properties, the two forms bias orienting in different ways, by operating at different stages of processing between stimulation and response. As we will show, one form does so by biasing perception, whereas the other form does so by biasing action. We will refer to the inhibition that Posner and Cohen described as an “input” form and the form of inhibition that Posner and colleagues described as an “output” form. The possible loci of operation of these two forms of IOR are illustrated by explicitly placing them into a theoretical framework that clearly distinguishes input from output stages of processing (for example, see Figure 20.1). The input form of IOR operates relatively early in the processing of stimuli, perhaps by reducing the salience of stimuli near the inhibited object or location. The output form operates relatively late in processing as a bias against spatial responses (either manual or oculomotor) that may include covert orienting.

## Diagnostics for Distinguishing between the Two Forms of Inhibition

In the remainder of this chapter we will use two diagnostics to show that the input effect is generated when the oculomotor system responsible for reflexive saccades is suppressed, whereas the output effect is generated when this reflexive oculomotor system is *not* suppressed. By employing this two-form conceptualization of IOR we are also able to reconcile some divergent findings from the literature, including some from within our own work.

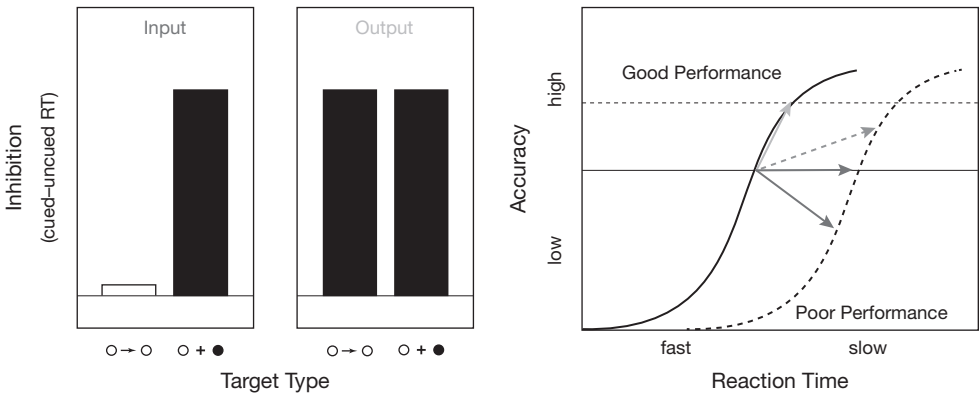


**Figure 20.1** A figure illustrating the flow of information between cognitive modules from perceptual (P), input-related processing to motoric (M), output-related processing. A central stage of processing (H), which can be bypassed, is the site of awareness of perceptually derived categories and objects and of the initiation of some actions. The two dashed arrows have been added by RMK and are not in the original figure, which has been redrawn from Attneave (1961). We have inserted into this framework (red and green surfaces) where the two forms of inhibition described in this chapter might be operating. A black and white version of this figure will appear in some formats. For the colour version, please refer to the plate section.

## Central vs. Peripheral Arrow Targets

One diagnostic for these two forms of inhibition depends on comparing performance with peripheral onset vs. central arrow targets, with both types of target calling for the same left–right localization response either manually or with a saccadic response, usually in the direction of the peripheral stimulus or the direction indicated by the central arrow. First employed independently by Abrams and Dobkin (1994) and Rafal, Egly, and Rhodes (1994), the logic behind this diagnostic is as follows: If the inhibition exclusively delays or degrades the flow of information from the region of space in which it has been generated, then, while responses to cued peripheral targets should show a RT delay relative to uncued peripheral targets, responses to centrally presented arrows should not be affected by the cue (Figure 20.2A, left panel). On the other hand, if the inhibition exclusively biases responding, then the two types of target should show equal amounts of RT delay in the cued direction (Figure 20.2A, right panel).

When this diagnostic was used by Tracy Taylor in her 1997 dissertation (later described in Taylor & Klein, 2000), she discovered a dramatic dissociation. Six “experiments” were generated from the orthogonal combinations of responses to the first (S1: ignore, manual, saccade) and second (S2: manual, saccade) stimuli, which were separated by a 1 second CTOA and a cue back. In a block of trials, both S1 and S2 were, randomly, either peripheral events or central arrows. S1s were uninformative about

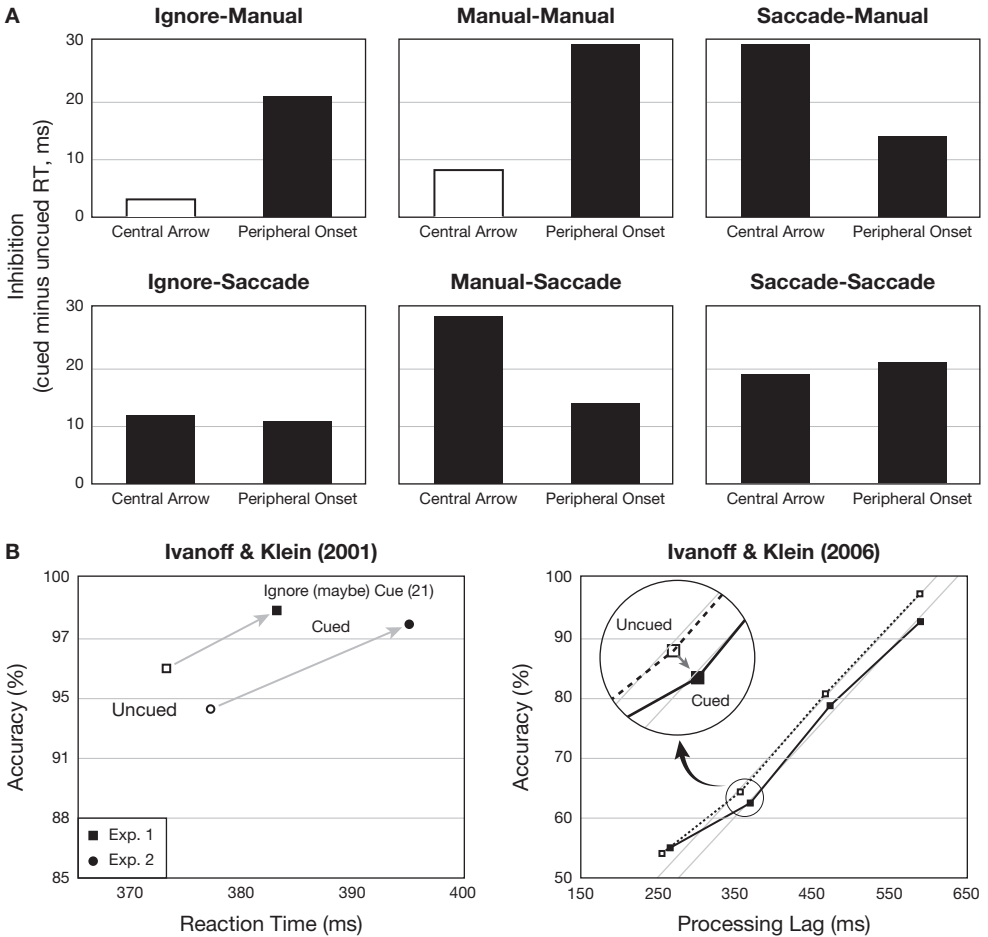


**Figure 20.2** Diagnostics for determining whether an effect is early (input) or late (decision/output) in processing. (A) Target-type diagnostic. Two hypothetical patterns of results when the different target types (central arrow or peripheral target, shown on the x-axis) are calling for a localization response are presented about 1 second after an uninformative peripheral cue. (B) Speed-accuracy diagnostic. A nonspatial, two-alternative, forced-choice response is required to a continuous target, and performance is examined in speed-accuracy space. Two hypothetical information accrual functions (solid and dashed curves) and two hypothetical criteria for responding (horizontal lines) are illustrated. See text for explanation. A black and white version of this figure will appear in some formats. For the colour version, please refer to the plate section.

the upcoming S2s. Although the six combinations of what participants were instructed to do in response to S1 and S2 – when orthogonally combined with the two stimulus types of each S1 and S2 – yielded 24 unique conditions for which IOR could be measured, here we will focus only on the conditions with peripheral S1s. The results from these 12 conditions are illustrated in Figure 20.3A, where a significant inhibitory effect is represented by filled (black) bars and the absence of inhibition is represented by unfilled (white) bars. Importantly, two dramatically different forms of inhibition were apparent. The *input* form, observed when there was no eye movement to either S1 or S2, is characterized by inhibition in response to peripheral onset targets but not in response to central arrow targets (as in the left panel of Figure 20.2A). The *output* form, observed when a saccade was made to either S1 or S2 (or both), is characterized by significant, and roughly equivalent amounts of inhibition in response to both peripheral onset and to central arrow targets (as in the right panel of Figure 20.2A).

## Examining Results in Speed–Accuracy Space

A second diagnostic was first implied by Ivanoff, Klein, and Lupianez (2002) in their analysis of the interaction between the inhibitory aftereffect of a peripheral cue and the Simon effect. Representing the temporal dynamics of information processing in speed-accuracy space (see Figure 20.2B), Ivanoff and colleagues noted two different mechanisms by which a hypothetical inhibition might cause a delay in RT. Without affecting the accumulation of information about a continuously presented target, one inhibitory effect might raise the threshold or criterion for responding. When such a “response bias” form of inhibition was operating, the change in performance would be characterized as a trade-off between speed and accuracy: that is, performance for cued targets would be slower *but also* more accurate than for uncued targets (a shift along a single SAT function as illustrated by the green arrow in Figure 20.2B). In contrast, the other form of inhibition



**Figure 20.3** (A) Results from the target-type diagnostic as employed in Taylor’s dissertation (Taylor & Klein, 2000). Significant IOR is represented by black bars; non-significant IOR is represented by white bars. Each panel represents a different combination of tasks the participant performed in response to S1 and S2 (e.g., Ignore-Manual denotes ignoring S1 and making a manual response to S2). (B) Results from the speed–accuracy diagnostic as employed in Ivanoff’s dissertation (Ivanoff & Klein, 2001, 2006). RT and accuracy are plotted on the x- and y-axes, respectively, for targets presented at the cued (filled symbols) and uncued (unfilled symbols) conditions. In the left panel participants performed a go/no-go task with typical reaction time methodology. In the right panel a response-window technique was used to generate full SAT functions. The green arrows signify that IOR is characterized by a speed–accuracy trade-off and the possibility of an output effect. The red arrow signifies a genuine decrease in information processing efficiency. A black and white version of this figure will appear in some formats. For the colour version, please refer to the plate section.

delays or degrades information processing about either the target or the connection between the target and the correct response.

When an input form of inhibition is operating, as represented by the red arrows in Figure 20.2B, performance in response to cued targets should fall on a different and less efficient speed–accuracy trade-off (SAT) function. The inhibited function could be shifted rightward from the performance with uncued targets, as illustrated in Figure 20.2B, or it might begin to rise above chance at the same time while increasing to asymptote more

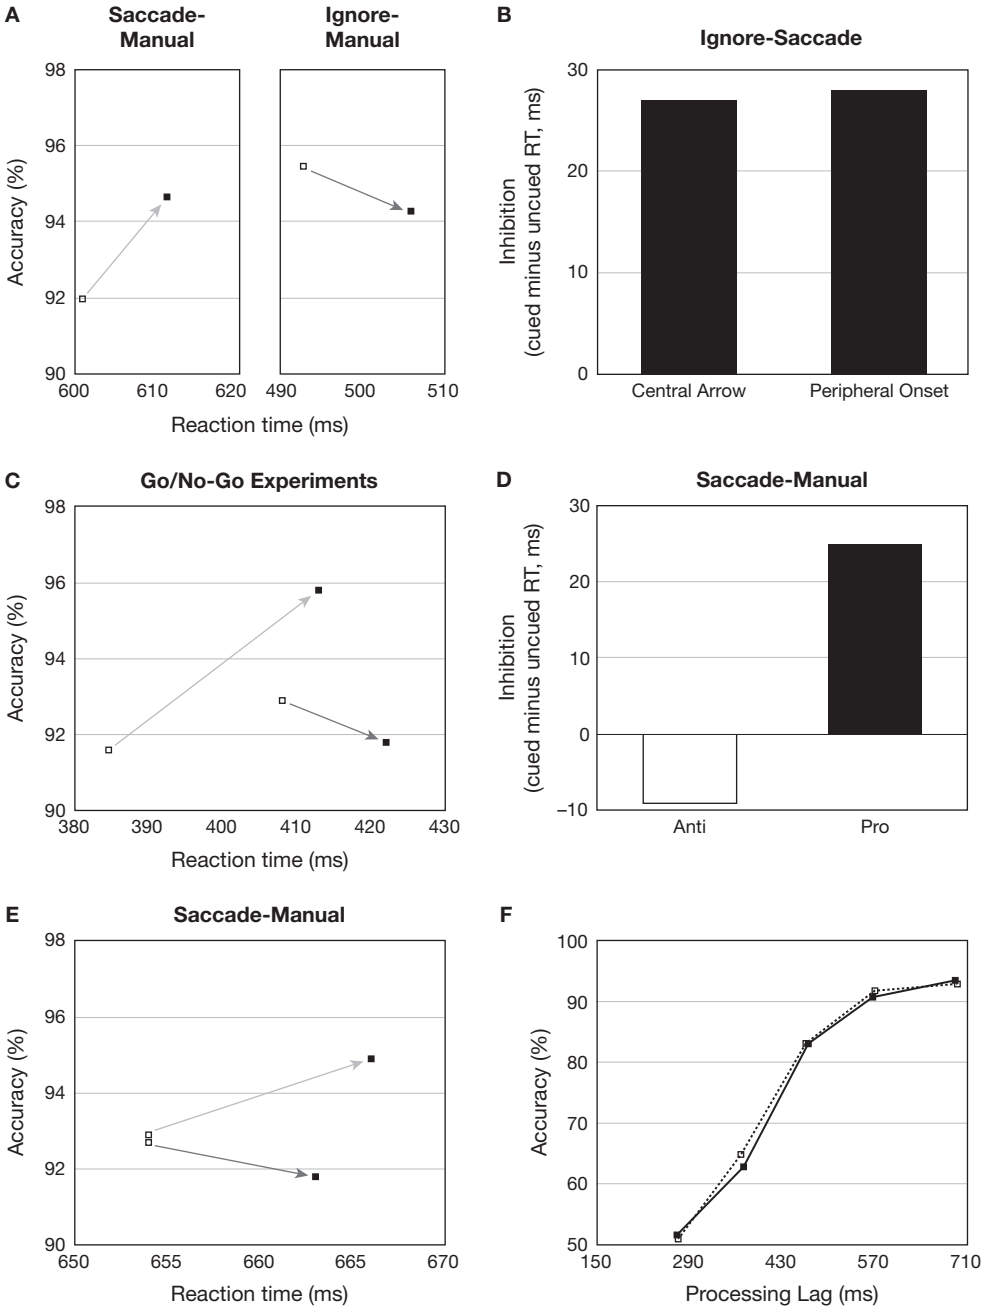
slowly than for uncued targets (not illustrated). RTs to targets suffering from this form of inhibition will be slower than to uncued targets and accompanied by either no change in accuracy or worse accuracy.

This speed–accuracy space diagnostic requires continuously presented peripheral targets calling for a nonspatial discrimination (such as a go/no-go task or two alternative forced choice [2-AFC] task about a property that is orthogonal to the cuing dimension) about the form, identity, or color of a target. The resulting pattern of performance is considered in speed–accuracy space. Using such an RT task, these two different effects on target processing can be distinguished by the direction (or slope) of the line connecting performance (speed and accuracy) from uncued to cued (inhibited) targets. When this line is horizontal or negatively sloped (red lines in Figure 20.2B), there is a genuine effect on the quality of target processing (an RT cost without an increase in accuracy or a cost in both RT and accuracy). When the effect is due to a bias (a change in the criterion for responding), this line will have a positive slope (green line in Figure 20.2B). Whereas the flat or negatively sloping relationship between cued and uncued performance in SAT space is relatively unambiguous in indicating that a shift in the SAT function has been obtained, unfortunately a positively sloping relationship is ambiguous: It can be due to a response bias (green line) or it could be due to a combination of bias and degraded information processing (blue dashed line).

An even more powerful method (but a much more costly one) entails generating full SAT functions (a methodology advocated by Wicklegren, 1977) for cued and uncued targets. The SAT method allows visualization of the temporal dynamics of information processing for cued and uncued targets. And, when the input form of inhibition is being observed, this method allows the scientist to determine which parameter (e.g., start time, rate, asymptote) is affected by the inhibition (cuing). Another benefit of this more powerful method is that it can be used to disambiguate the situation described earlier; is a positive slope coming from a single function or from different functions? Both of these methods were pioneered by Jason Ivanoff in a series of experiments beginning with his 2003 dissertation (Hilchey, Ivanoff, Taylor, & Klein, 2011; Ivanoff & Klein, 2001, 2004, 2006; Ivanoff, Klein, & Lupianez, 2002).

Ivanoff and Klein (2001) explored inhibitory aftereffects of peripheral cues using a go/no-go task. In prior studies using the simple detection task with catch trials, false alarms could be observed, but because nothing was presented on a catch trial, they could not be attributed to the cued versus uncued location. In a go/no-go task, stimuli are presented on every trial at the cued or uncued locations. One stimulus (go) calls for a response; the other one (no-go) does not. Using such a go/no-go task allowed Ivanoff to be able to attribute false alarm rates to the cued vs. uncued locations. In this study (see Figure 20.3B, left panel), he discovered a pattern that was characteristically a response bias: slower responding to cued targets was accompanied by higher accuracy (there were fewer false alarms for cued targets). In two experiments from his dissertation, Ivanoff (presented in Ivanoff & Klein, 2006) used the SAT method (cf. Wicklegren, 1977) to generate full SAT functions for targets presented at cued and uncued locations. Here (see Figures 20.3B, right panel), he demonstrated the input form of inhibitory aftereffect: performance at the cued location was shifted rightward in SAT space.

A few years later it occurred to Klein that if the effect of IOR generated by an eye movement was, as proposed by Posner, Rafal, Choate, and Vaughan (1985) and confirmed by Taylor and Klein (2000), just a bias against responding in the cued direction,



**Figure 20.4** Patterns of findings when the two diagnostics illustrated in Figure 20.2 are applied to our studies and to the literature. In A, C, E, and F, data from cued targets are plotted with filled symbols; data from uncued targets with unfilled symbols. (A) Results from Chica et al. (2010). Data connected by the green arrow are from Experiments 1B and 3B; data connected by the red arrow are from 3A. (B) Data from Hilchey, Klein, & Ivanoff (2012). (C) Data are from a go/no-go tasks (errors are false alarms). Data connected by the red arrow are from Hilchey et al. (2014) in which there was careful monitoring of eye position and participants were given feedback whenever an untoward eye movement occurred. Data connected by the green arrow are collapsed from six studies in



then IOR might not be observed when caused by an eye movement and measured with a nonspatial discrimination task. It might seem counterintuitive to propose that something would be caused with no apparent effect. In Klein’s thinking, at that time, the hypothetical response bias would have been generated by the eye movement to the signal in the periphery, but it would not have been measurable by the 2-AFC task, because targets at the cued location did not require spatial (localization or detection) responses. Klein suggested this to Ana Chica who was visiting his laboratory from the University of Granada, and Chica, Taylor, Lupianez and Klein (2010) conducted just such an experiment. The results provided a dramatic link between the ideas of Klein’s former students, Taylor and Ivanoff.

Chica, Taylor, Lupianez, and Klein (2010) explored two combinations from the matrix illustrated in Figure 20.3A: S1 was either ignored or a saccade target, and S2, presented in the periphery (at the cued or opposite location), required a nonspatial, 2-AFC discrimination based on color. In both cases, eye position was monitored. We expected degraded processing along the input pathway following ignored cues and no such degradation – but possibly a response bias – when saccades were made to the cue. Confirming our predictions, the results (see Figure 20.4A) showed that when participants were instructed to ignore the cue, maintain fixation, and were “beeped” if they made an eye movement, the inhibition was of the input form (right panel). In contrast, when subjects successfully made a saccade to the cue and back to the original fixation before the final target (left panel), the output form (a speed–accuracy trade-off) was observed.

## Using the Diagnostics to Resolve Puzzles in the Literature

We have since explored three puzzles concerning the pattern of results from some of conditions illustrated in Figure 20.3A.

### Ignore–Saccade: Input + Output or Pure Output?

One puzzle is why Taylor and Klein found equivalent inhibitory effects when saccades were made to peripheral and central arrow targets after an ignored peripheral cue (a pure output form as illustrated in Figure 20.2A, right panel), whereas Abrams and Dobkin (1994) found greater IOR following peripheral than central targets (both input and output forms of IOR). Abrams and Dobkin inferred from this pattern that two components were jointly contributing to the inhibition in this condition: an output component that delayed responses in the cued direction (and that would be operating with both types of target) and an input component that affected only the peripheral target. Klein and Hilchey (2011, p. 480) proposed the following resolution to this conflicting pattern of results:

---

**Figure 20.4** (cont.) the literature (Ivanoff & Klein, 2001, 2003, 2004, Exp. 1; Taylor & Ivanoff, 2003; Prime & Ward, 2008, Exp. 3; Prime & Jolicoeur, 2009, Exp. 1). (D) Data are from Hilchey, Dohmen, Crowder, & Klein (2016). In separate blocks participants made pro- or anti-saccades in response to a peripheral luminance increment, and ~1 second after the onset of this stimulus they made a manual response in the direction of a centrally presented arrow. (E) Data are from Redden, Hilchey, & Klein (2016). Different participants made a pro- or anti-saccade in response to a peripheral luminance increment, and ~1 second after the onset of this stimulus they made a nonspatial 2-AFC. (F) From Redden, Hilchey, and Klein (2015a). Participants made a pro-saccade in response to a peripheral luminance increment, and ~1 second after the onset of this stimulus a target was presented calling for a nonspatial 2-AFC response. Importantly, speed–accuracy trade-off functions were obtained for cued and uncued targets using a window method. A black and white version of this figure will appear in some formats. For the colour version, please refer to the plate section.

Abrams and Dobkin's finding was generated by untoward differences in mental set made possible by their collecting data from arrow and peripheral targets in separate blocks. When the target was never in the periphery (always an arrow at fixation), the participant could adopt an attentional control setting that encouraged effective filtering of the irrelevant peripheral cues. This would not be possible in a block when all (Abrams and Dobkin's peripheral target blocks) or some (Taylor and Klein, 2000) of the targets were peripheral.

Hilchey, Klein, and Ivanoff (2012) confirmed that the Abrams and Dobkin pattern was compromised by an untoward design feature: The two target types were presented in separate blocks of trials. When the two target types are appropriately randomly intermixed, the Taylor and Klein finding of roughly equivalent IOR for the two target types is confirmed (Figure 20.4B).

## No Response—Manual: Input or Output?

Another puzzle concerns the fact that in the Ignore-Manual combination with peripheral cues and targets, a speed–accuracy trade-off has often been reported in a go/no-go task even when participants are instructed not to move their eyes. This pattern, first reported by Ivanoff and Klein (2001; see Figure 20.3B, left panel), has been observed in five subsequent studies (see green arrow in Figure 20.4C). We surmised that despite instructions, perhaps the reflexive oculomotor system was not sufficiently inhibited in these studies for the input form to have been generated, because in none of these studies were participants given trial-by-trial feedback about any unwanted eye movements.

We (Hilchey et al., 2014) therefore closely replicated the methods of Ivanoff and Klein (2001) while monitoring eye movements and, importantly, providing trial-by-trial feedback whenever an unwanted eye movement occurred. Under these conditions (see the red arrow in Figure 20.4C), cued targets showed the typical, significant, RT delay, but there was no evidence for a speed–accuracy trade-off; that is, responses to cued targets were not more accurate than those to uncued targets. It is worth mentioning that we found one experiment (Cheal, Chastain, & Lyon, 1998) in which a go/no-go task was combined with assiduous monitoring of eye position and feedback when untoward eye movements were made. In this experiment, there was an RT delay for cued targets and no evidence of a speed–accuracy trade-off. Our pattern of results when considered together with the literature shows the importance of not only monitoring eye position but also giving participants trial-by-trial feedback when unwanted eye movements are executed.

## Saccade—Manual: Output or Input?

A third puzzle arises from studies that have asserted that an input form of IOR is generated by saccadic eye movements, whereas, in contrast, Taylor's findings imply an output form. The challenge is rooted in the finding that following an anti-saccade (a saccade executed away from a peripheral "target"), the inhibitory aftereffect is at the location of the stimulus and not the response (e.g., Fecteau, Au, Armstrong, & Munoz, 2004). To resolve the conflict, Klein and Hilchey (2011) proposed that the input form of IOR ought to be generated by an anti-saccade, because in order to successfully look away from a peripheral target, the reflexive oculomotor machinery must be tonically inhibited (Forbes & Klein, 1996; Ignaschenkova et al., 2004). This view, which is critical for the present chapter, represents an important qualification upon Taylor's inference that the output form of IOR is generated whenever eye movements were made.

We (Hilchey, Dohmen, Crowder, & Klein, 2016) first employed the central-target diagnostic in two conditions in which an eye movement was made in response to a peripheral stimulus and then, upon returning gaze to fixation, a manual response was required to respond to a central arrow. Replicating the methods of Taylor and Klein (2000), in one condition, the eye movements were pro-saccades. Differing from Taylor and Klein, however, there were no arrow S1s or peripheral target S2s. If Taylor and Klein’s finding of IOR in this condition were due to the fact that their participants experienced both peripheral and central stimuli as both S1s and S2s and if the form of IOR generated by saccades to peripheral targets really is input-based, then no IOR should be observed in this condition. In contrast, if the output form of IOR is generated when saccades are executed to peripheral targets, then we should replicate Taylor and Klein’s finding of IOR when it is measured using manual responses to central arrow targets. In a second condition we required participants to make an anti-saccade to the peripheral S1. If Klein and Hilchey’s (2011) proposal, that an input form of IOR would be generated in this condition because the reflexive oculomotor machinery must be tonically inhibited for correct anti-saccade performance, were correct, then there would be no IOR observed in manual responses to central arrows. As seen in the “pro-saccade” bar of Figure 20.4D, 22 msec of significant IOR was observed. Replicating Taylor and Klein’s finding when this condition was mixed with others, this finding confirms that pro-saccades generate the output form of IOR. And, consistent with Klein and Hilchey’s proposal that the input form would be generated when an anti-saccade was used to cause the inhibitory effect, manual responses to centrally presented arrows were not delayed in the cued direction (Figure 20.4D, anti-saccade bar).

Converging evidence for this interpretation is provided by Redden, Hilchey, and Klein (2016) who applied the SAT diagnostic to this cell. As in Hilchey et al. (2016), IOR was generated by a pro- or anti- saccade to a peripheral target. But following the SAT diagnostic, we measured the effect of these two inhibitions in a nonspatial discrimination task. Replicating Chica and colleagues, we found a speed–accuracy trade-off (Figure 20.4E, green arrow) when IOR was generated by a pro-saccade. Reinforcing our interpretation of the results in Figure 20.4D, when generated by an anti-saccade, an input form of inhibition was observed (Figure 20.4E, red arrow). Next, using the SAT method to measure accuracy while controlling RT, Redden, Hilchey and Klein (2015a) generated SAT functions for cued and uncued targets in conditions otherwise the same as those used here with pro-saccades. As would be predicted by an output-based account, when IOR was caused by a pro-saccade, cuing had no effect on the SAT functions (Figure 20.4F). Compare this pattern with SAT functions from the ignore-manual cell of this matrix (Figure 20.3B, right panel), in which Ivanoff and Klein (2006) found that IOR following an ignored cue caused a significant rightward shift of the SAT function, which implies a degradation or delay in the processing of cued targets in the ignore-manual condition of this matrix.

## Summary

Using the diagnostic patterns obtained when peripheral onset or central arrow targets are used, we have seen that there are two forms of inhibitory aftereffect. These are distinguished clearly by whether the effect can only be measured by peripheral targets (the input form when the reflexive oculomotor system is suppressed) or by whether there are roughly equivalent delays in response whether the targets are central or peripheral (the output form

when the reflexive oculomotor system is not suppressed). When performance on cued and uncued targets is represented in SAT space, the input form is manifest as a shift in the SAT function representing degraded or delayed processing of cued targets, whereas the output form entails NO shift in the function, but instead a movement along it (a response bias).

It has also been shown that the input form of IOR is represented in object-based coordinates (Hilchey, Pratt, & Christie, 2016), whereas the output form of IOR is represented in spatiotopic coordinates (Christie, Hilchey, & Klein, 2013; Redden, Hilchey, & Klein, 2015b), suggesting yet another dissociation at the neural level. This “object-” vs. “space-” representation suggests the input form is likely modulated cortically in the ventral visual stream, whereas the output form is likely modulated by the interaction between subcortical structures (vis superior colliculus) and cortical areas in the dorsal visual stream.

In conclusion, and speaking cognitively, either form of inhibition can perform the function attributed to IOR by Posner – “to maximize sampling of the visual environment.” Using Attneave’s thoughtful representation of a cognitive architecture (Figure 20.1), input attenuation operates on a salience map that influences what we will attend to and become aware of while a response bias operates on a priority map that influences what behaviors (including orienting) we are likely to engage in.