

Reduced Effectiveness of Interruptive Drug-Drug Interaction Alerts after Conversion to a Commercial Electronic Health Record

Adam Wright, PhD^{1,2,3}, Skye Aaron, BA¹, Diane L. Seger, RPh^{1,3}, Lipika Samal, MD, MPH^{1,2,3}, Gordon D. Schiff, MD^{1,2,3}, and David W. Bates, MD, MSc^{1,2,3}

¹Division of General Internal Medicine and Primary Care, Brigham and Women's Hospital and Harvard Medical School, Boston, MA, USA;

²Department of Biomedical Informatics, Harvard Medical School, Boston, MA, USA; ³Information Systems Department, Partners HealthCare, Boston, MA, USA.

BACKGROUND: Drug-drug interaction (DDI) alerts in electronic health records (EHRs) can help prevent adverse drug events, but such alerts are frequently overridden, raising concerns about their clinical usefulness and contribution to alert fatigue.

OBJECTIVE: To study the effect of conversion to a commercial EHR on DDI alert and acceptance rates.

DESIGN: Two before-and-after studies.

PARTICIPANTS: 3277 clinicians who received a DDI alert in the outpatient setting.

INTERVENTION: Introduction of a new, commercial EHR and subsequent adjustment of DDI alerting criteria.

MAIN MEASURES: Alert burden and proportion of alerts accepted.

KEY RESULTS: Overall interruptive DDI alert burden increased by a factor of 6 from the legacy EHR to the commercial EHR. The acceptance rate for the most severe alerts fell from 100 to 8.4%, and from 29.3 to 7.5% for medium severity alerts ($P < 0.001$). After disabling the least severe alerts, total DDI alert burden fell by 50.5%, and acceptance of Tier 1 alerts rose from 9.1 to 12.7% ($P < 0.01$).

CONCLUSIONS: Changing from a highly tailored DDI alerting system to a more general one as part of an EHR conversion decreased acceptance of DDI alerts and increased alert burden on users. The decrease in acceptance rates cannot be fully explained by differences in the clinical knowledge base, nor can it be fully explained by alert fatigue associated with increased alert burden. Instead, workflow factors probably predominate, including timing of alerts in the prescribing process, lack of differentiation of more and less severe alerts, and features of how users interact with alerts.

KEY WORDS: clinical decision support; adverse drug events; drug-drug interactions; electronic health records; safety.

J Gen Intern Med 33(11):1868–76

DOI: 10.1007/s11606-018-4415-9

© Society of General Internal Medicine 2018

Received July 13, 2017

Revised December 1, 2017

Accepted March 16, 2018

Published online May 15, 2018

BACKGROUND

Adverse drug events (ADEs) represent a major cause of harm in healthcare.^{1–4} Many ADEs are preventable,^{5–7} and computerized provider order entry (CPOE) with clinical decision support (CDS) reduces the frequency of ADEs.^{8,9} Preventable ADEs have myriad causes, including prescribing medications that a patient is allergic to, prescribing nephrotoxic drugs in the setting of renal dysfunction, making dosing errors, or prescribing multiple interacting drugs.⁶ Numerous CDS systems have been implemented to reduce harm from each of these causes, but the most widely used type of CDS is drug-drug interaction (DDI) checking.^{10,11} DDI checking has reduced the rate of ADEs and is a critical tool to support safe prescribing in electronic health records (EHRs).^{9,12–14} Use of DDI checking is also a certification and attestation requirement in all three stages of the United States Meaningful Use incentive program for EHRs.¹⁵

A DDI is considered present when one drug alters the effects of another. Some DDIs cause ADEs that would not have occurred if either of the drugs had been prescribed by itself. Mechanisms for drug interactions may include pharmacokinetic interactions, where “one drug affects the absorption, distribution, metabolism, or excretion of another” and pharmacodynamic interactions, where “two drugs have additive or antagonistic pharmacologic effects.”¹⁶ Automated DDI checking in EHRs checks medication orders against a list of known DDIs and alerts clinicians when they prescribe two drugs together that are on the list of known DDIs. At the extremes—the most severe and least severe DDIs—there is consensus about which are important and unimportant.^{17,18} However, for many of the DDIs in between, there is no broad consensus about their level of importance or which ones merit alerts. Various commercial knowledge bases cataloging potential DDIs are available, but there is little agreement about the drug pairs to include, and there is a paucity of primary literature supporting many purported interactions.^{18,19}

Consequently, over the last two decades, Brigham and Women's Hospital (BWH) developed a highly targeted knowledge base and alerting system for DDIs focusing on key interactions and minimizing the burden of over-alerting.²⁰ The customized and highly tailored knowledge base was

designed by a team of pharmacists and physicians.²¹ Compared to most commercially available drug knowledge bases as typically installed, the homegrown BWH knowledge base triggers alerts much less frequently. The alerting system, which was built into BWH's legacy EHRs, the Longitudinal Medical Record (LMR) and the Brigham Integrated Computer System, strongly differentiates between various severity tiers for DDIs; for example, contraindicated drug pairs could not be ordered at all, while the least severe interactions were informational and did not interrupt workflow.²⁰

BWH recently adopted the Epic commercial EHR system (Verona, WI). In this study, we evaluated the effect of the transition from our internally developed LMR to our implementation of Epic on DDI alerting behavior during medication order entry, including alert burden, acceptance rate, and alert fatigue. All evaluation and discussion of Epic in this paper applies to the specific implementation of Epic at BWH, which may differ from Epic implementations at other institutions.

METHODS

BWH transitioned from LMR to Epic on May 31, 2015. As part of the transition, BWH switched from using a highly tailored DDI knowledge base to a widely used commercial knowledge base from First Databank produced by Hearst Corporation (South San Francisco, CA) with minimal customization. This is a frequently used approach—typically with the implementation of commercial EHRs; one of several commercial drug knowledge bases sold by third parties is selected and used.

Figure 1 shows the strongly tiered alerting system that was used in LMR, and the three tiers are described in Table 1. Figure 2 shows the weakly tiered alerting system that is used in the Epic implementation at BWH. Although both systems use three tiers of drugs, the drugs contained in each tier differ, and the difference in alerting between tiers is less strong in Epic than in LMR. In our implementation of Epic, all DDI alerts are interruptive regardless of tier and may be overridden if desired, with room for an optional override reason. Unlike in LMR (where users were prohibited from ordering the most severe drug combinations while the least severe combinations resulted in only a passive alert), Epic at BWH does not use hard stops (which make co-prescription of interacting drugs impossible) or passive alerts. Thus, the severity of tiers in Epic is differentiated only by color, icon, and label.

Additionally, placement of DDI alerts within provider workflow changed significantly with Epic (see Fig. 2). While DDIs were shown as soon as a medication was selected in LMR, in Epic as implemented at BWH, the DDI alerts are shown in one batch at the time of signing. Thus, in LMR, providers see the alert at the beginning of the ordering process, while in Epic, providers see the alert at the end of the ordering process. The result is that, in our implementation of Epic, users may see many medication alerts at once, since multiple orders

can be signed at the same time, and the DDI alerting checks all the pairs at once. Furthermore, in our implementation of Epic, users may override all alerts with a single click—there is no requirement to review or interact with each alert separately, as was required in LMR.

We extracted data on outpatient medication orders and alert firing and acceptance in LMR from our enterprise data warehouse from November 28, 2014, to May 29, 2015—the last 6 months that LMR was used at BWH. We also extracted outpatient Epic data from June 1, 2015, to November 30, 2015—the first 6 months that Epic was used. Although Epic is used in both the inpatient and outpatient setting, we limited our analysis to outpatient alerts because LMR was used only in the outpatient setting. BWH employed a “big bang” approach to the system implementation, converting all outpatient clinics and inpatient areas on the same day. Because we transitioned from LMR to Epic on a Sunday morning, we excluded the Saturday and Sunday surrounding the cutover.

We defined alert burden as the number of alerts per 100 orders, and we defined clinician acceptance rate as the proportion of DDI alerts leading to either the cancelation of the proposed drug order or discontinuation of the incumbent interacting drug. Note that we did not attempt to assess the appropriateness of the clinical response to the DDI alert. For example, while many less severe DDI drug pairs may be safely prescribed so long as the patient is monitored, we did not attempt to include monitoring in our definition of alert acceptance. Additionally, we did not attempt to control for instances where a user accepted or overrode an alert but, in practice did something else, either by continuing to prescribe a drug after accepting an alert or temporarily discontinuing a drug after overriding an alert. We looked only at the action the user took on the DDI alert itself to evaluate the impact of the two different alerting systems.

We compared the alert burden and the alert acceptance rates in the last 6 months of LMR use to the first 6 months of Epic use. We compared alerts overall, by tier, and for the most commonly alerted drug pairs in each system. As described in the “Results” section, our leadership disabled Tier 3 alerting in Epic on March 1, 2016, so we were also able to further analyze alert fatigue by comparing alert burden and acceptance rates before and after this change. This analysis was done using a separate time period, using August 31, 2015, through February 29, 2016, as the “before” period, and March 2, 2016, through August 31, 2016, as the “after” period.

RESULTS

As shown in Figure 3, the number of interruptive DDI alerts (which included Tiers 1 and 2 in LMR, and all three tiers during the initial period in Epic) increased substantially after the switch to Epic, while total volume of medication orders stayed about the same. Because each system logs edits to medication orders and refills differently, the medication

volume is similar but not exactly comparable between systems. It does however suggest that prescription volumes were not substantially different after the switch from LMR to Epic. In the final 6 months of LMR, there were approximately 2 interruptive DDI alerts per 100 orders, while in the first 6 months of Epic implementation, there were about 11 interruptive DDI alerts per 100 orders. Both systems also display many other types of medication alerts, including warnings for high and low doses, drug and therapeutic duplication, potential drug-allergy issues and pregnancy, geriatric and pediatric precautions. In LMR, DDI alerts comprised 19.2% of all medication alerts, leading to a total alert burden of about 9 medication alerts per 100 medication orders. In our instance of

Epic, DDI alerts comprised about 29.8% of all medication alerts, leading to a total alert burden of about 36 medication alerts per 100 medication orders.

The acceptance rate for interruptive DDI alerts overall (Tiers 1 and 2 in LMR and Tiers 1, 2, and 3 in Epic) fell from 30.7% in LMR to 6.2% in Epic ($P < 0.001$ using Pearson's chi-squared test for independence). A comparison is not possible for Tier 3 because these alerts were passive in LMR. Comparing only Tiers 1 and 2 in each system, the overall acceptance rate fell from 30.7% in LMR to 7.6% in Epic ($P < 0.001$). As shown in Figure 4, the acceptance rate fell for DDIs in both Tiers 1 and 2. One possible explanation for this substantial drop in DDI acceptance is that many less

(A) Tier 1 (most severe)

(B) Tier 2 (moderate)


(C) Tier 3 (least severe)

Fig. 1 Screenshots of DDI warnings in LMR show. (A) Tier 1 (most severe) alert for azathioprine and febuxostat. User workflow is interrupted, and user must cancel one of the drugs. (B) Tier 2 (moderate) alert for citalopram and phenelzine. User workflow is interrupted, and user is given the option to cancel one of the drugs. (C) Tier 3 (least severe) alert for citalopram and ondansetron. Text in red passively notifies user of potential interaction without interrupting user workflow.

Table 1 Three Tiers of Interaction Severity Are Used in LMR

Tier	Severity	Alerting strategy	Examples
1	Most severe interactions; contraindicated	A red alert box is shown to the prescriber, who cannot proceed without discontinuing one or the other drug—this is a hard stop, and co-prescribing is not possible.	<ul style="list-style-type: none"> Phenelzine and dextroamphetamine (together these drugs may cause hypertensive crisis) Febuxostat and azathioprine (febuxostat can increase the effect of azathioprine, leading to bone marrow suppression²²)
2	Medium severity, but still potentially serious interactions	An interruptive yellow alert box is shown to the prescriber, who must provide a reason before being allowed to proceed; reasons include “Will adjust dose as recommended,” “Will monitor as recommended,” “Patient has already tolerated combination,” and “No reasonable alternatives.”	<ul style="list-style-type: none"> Clarithromycin and digoxin (clarithromycin can cause digoxin toxicity) Sulfamethoxazole/trimethoprim and warfarin (sulfamethoxazole/trimethoprim inhibits warfarin clearance which can lead to supratherapeutic anticoagulation²³)
3	Lower severity interactions	Rather than displaying an interruptive “pop up” alert, information about the potential interaction is shown at the top of the ordering screen, and no response from the user is required (or possible).	<ul style="list-style-type: none"> Omeprazole and voriconazole (may lead to increased levels and toxicity of voriconazole²⁴) Warfarin and simvastatin (may increase INR²⁵)

A Tier 1 (most severe)



B Tiers 2 (moderate) and 3 (least severe)

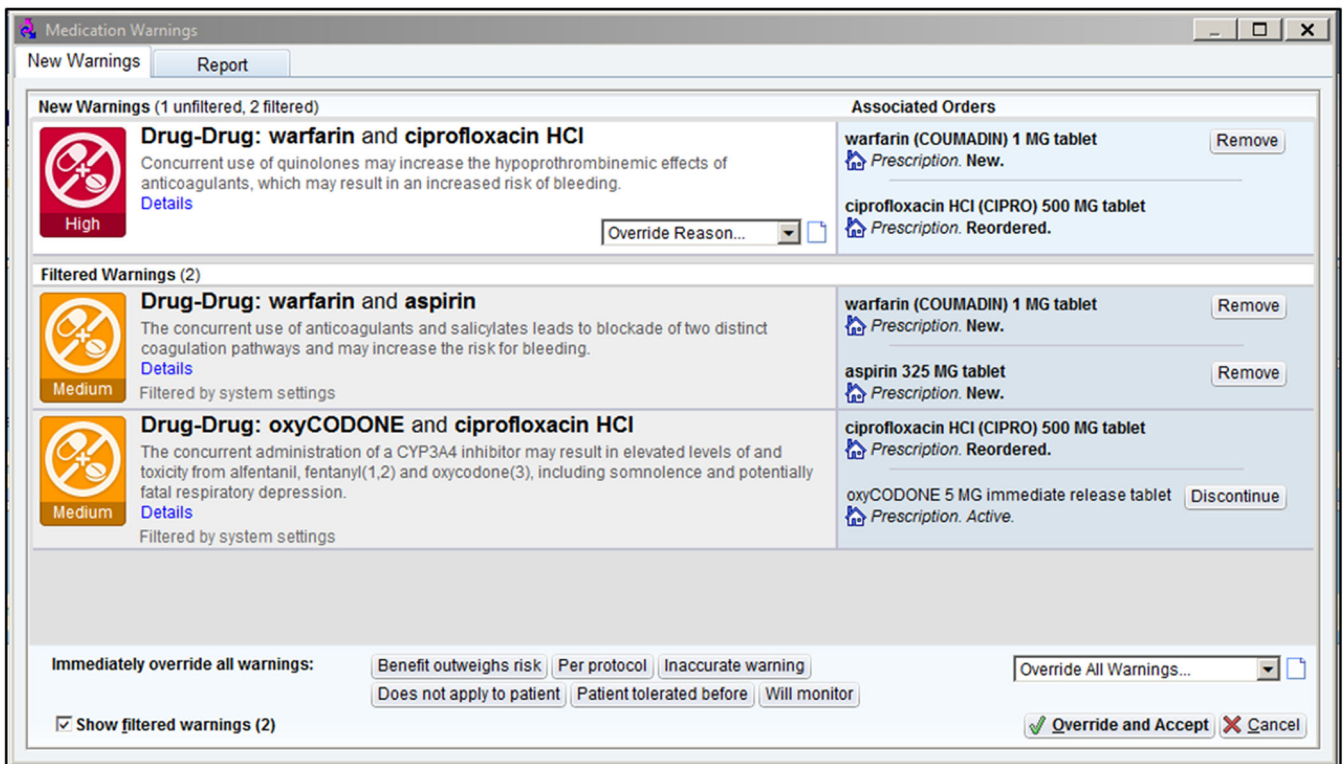


Fig. 2 Screenshots of DDI warnings in Epic show. (A) Tier 1 (most severe) alert for sumatriptan and phenelzine. User workflow is interrupted, and users are given the option of canceling one of the drugs. (B) Tier 2 (moderate) alert for warfarin and ciprofloxacin, and two Tier 3 (least severe) alerts for aspirin and warfarin, and ciprofloxacin and oxycodone. In both cases, user workflow is interrupted, and users are given the option of canceling one of the drugs. Due to concerns about excessive alerting and overrides (discussed below), Tier 3 alerts were altered on March 1, 2016. Before March 1, 2016, Tier 3 alerts were automatically displayed. After March 1, 2016, Tier 3 alerts are displayed to the user only when the user selects a checkbox to view them.

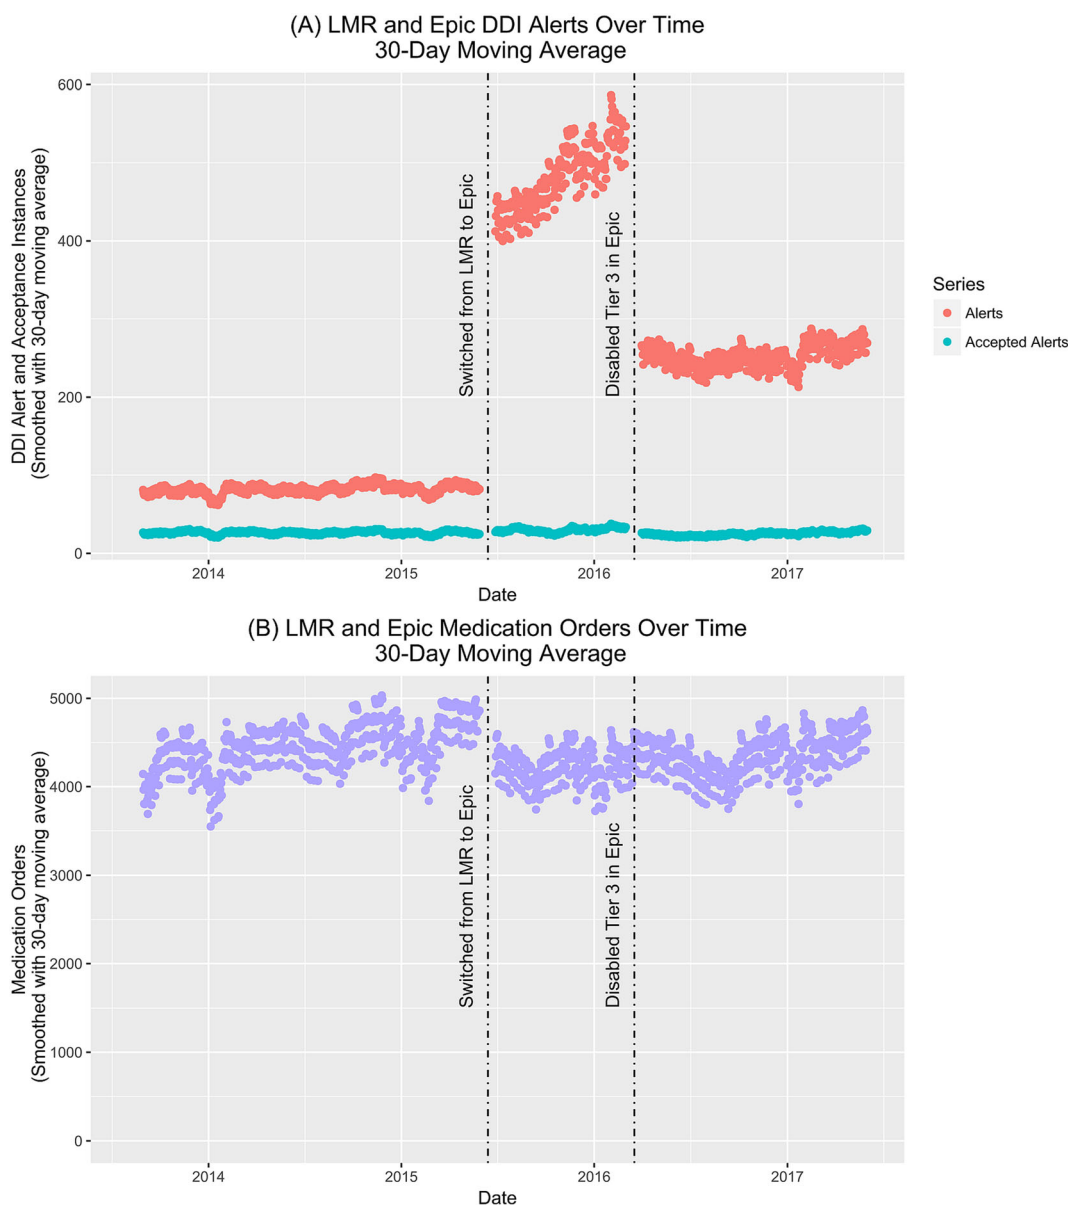


Fig. 3 (A) Number of DDI alerts and accepted DDI alerts over time, before and after the EHR transition, and after Tier 3 alerts were disabled. (B) Number of medication orders over time, before and after the EHR transition. Medication orders are logged differently in the two systems, but total volume per day is comparable.

important DDI alerts are shown by Epic and that these alerts are more likely to be overridden. This is certainly true, given that the knowledge base paired with Epic is much larger than the knowledge base paired with LMR. However, if this were the only cause, then we would expect identical pairs of interacting drugs that are present in both databases to have similar override rates in Epic and LMR. To explore this explanation, we looked at key identical DDI pairs in both systems to compare acceptance rates (shown in Table 2).

The three most common Tier 1 (most severe) alerts that fired in LMR were for ciprofloxacin and tizanidine,²⁶ clarithromycin and simvastatin,^{27, 28} and erythromycin and simvastatin.²⁸ These alerts were accepted 100% of the time in

LMR by design, but only 21.6, 18.0, and 15.6% of the time, respectively, in Epic.

When we evaluated Tier 2 DDIs, which are presented more similarly in Epic and LMR (i.e., they are interruptive but can be overridden in both systems), we found significant differences between clinician responses in the two systems. The three most common Tier 2 (medium severity) DDIs in LMR were citalopram and omeprazole²⁹ which fell from 29.6% acceptance in LMR to 8.9% in Epic, amlodipine and simvastatin³⁰ which fell from 20.6 to 5.8%, and ciprofloxacin and warfarin³¹ which fell from 31.7 to 17.9%.

Given the substantial drop in acceptance rates for DDI alerts, the Partners HealthCare CDS Committee hypothesized that the increased alert burden, particularly of Tier 3 alerts,

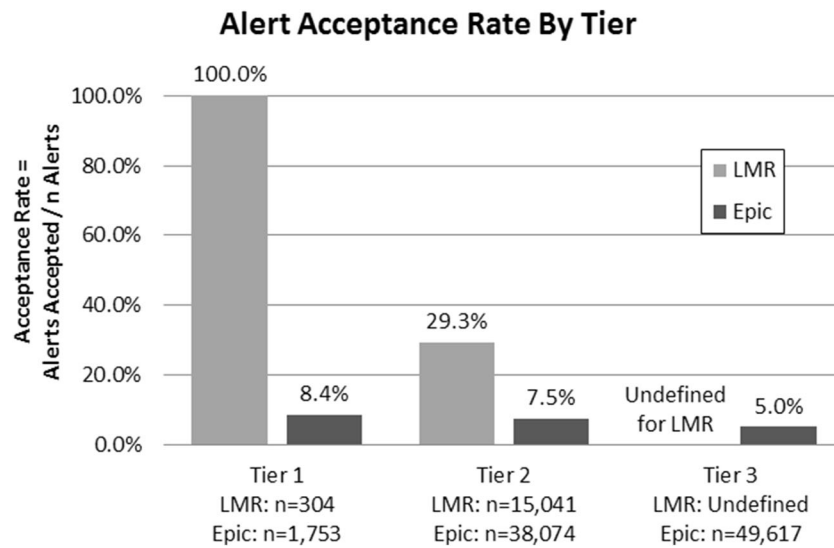


Fig. 4 Alert acceptance rates in LMR and Epic, by severity tier. Since the systems use different knowledge bases, the drugs contained in each tier differ according to the system. Acceptance rates are given as a percentage with the denominator in the axis label. The difference between rates in Tier 1 is statistically significant structurally with $P=0.0$. The difference between acceptance rates in Tier 2 is statistically significant with $P<0.001$ using Pearson’s chi-squared test for independence.

contributed to alert fatigue and the lower acceptance of Tier 1 and 2 alerts. The Committee decided to filter Tier 3 alerts so they would not be displayed unless specifically requested by the user. This change was implemented on March 1, 2016, and represents a natural experiment in alert fatigue. After the change, the interruptive DDI alert burden in Epic fell by 50.5%, as shown in Figure 3. Including other types of medication alerts (such as allergy or pregnancy), the total medication alert burden decreased from about 35 medication alerts per 100 orders to about 28 medication alerts per 100 orders in the 6 months before and after Tier 3 DDI alerts were disabled. After the change, Tier 1 alert acceptance rose from 9.1 to 12.7%, while Tier 2 acceptance rose from 6.9 to 9.2%. Though modest, these differences were statistically significant in both

tiers using Pearson’s chi-squared test for independence ($P<0.01$ for Tier 1 and $P<0.001$ for Tier 2).

DISCUSSION

We found a sixfold increase in the interruptive DDI alert firing rate when converting from a legacy system with a highly tailored drug knowledge base to a commercial system with a standard commercially available knowledge base. Even after disabling Tier 3 alerts, there were still almost three times as many interruptive DDI alerts in Epic as in LMR. This increase in alert burden was associated with a much higher override rate by providers. The increase in overrides was seen across severity

Table 2 Acceptance Rates for the Most Commonly Occurring DDI Alerts in LMR Are Shown, Along with the Corresponding Acceptance Rates in Epic. Acceptance Rates for Individual Drug Pairs Have Decreased from LMR to Epic. ¹Tier 1 LMR Acceptance Is 100.0% Because Providers Are Prohibited from Prescribing Both Drugs Together

Drug pair	LMR			Epic		
	Tier	Acceptances/alerts	Acceptance rate ¹ (%)	Tier	Acceptances/alerts	Acceptance rate (%)
Ciprofloxacin-tizanidine	1	56/56	100.0	1	8/37	21.6
Clarithromycin-simvastatin	1	44/44	100.0	1	11/61	18.0
Erythromycin-simvastatin	1	44/44	100.0	1	5/32	15.6
Gemfibrozil-simvastatin	1	18/18	100.0	1	1/27	3.7
Eplerenone-spirolactone	1	10/10	100.0	N/A	N/A	N/A
Prevpac-simvastatin	1	10/10	100.0	1	3/16	18.8
Citalopram-omeprazole	2	212/716	29.6	2	158/1773	8.9
Amlodipine-simvastatin	2	125/608	20.6	2	139/2391	5.8
Ciprofloxacin-warfarin	2	133/419	31.7	2	61/341	17.9
Cyclobenzaprine-tramadol	2	99/277	35.7	2	69/567	12.2
Levofloxacin-warfarin	2	71/265	26.8	2	16/161	9.9
Sildenafil-tamsulosin	2	42/256	16.4	2	20/452	4.4

levels, but particularly affected the most serious interactions, representing an important safety concern. The override rate for even the most severe warnings was about nine out of ten.

Factors Potentially Contributing to the Drop in Acceptance Rate

There are several possible explanations for the drop in interruptive DDI alert acceptance rates after this transition. First, interruptive DDI alerts fired 5.5 times more frequently in our initial implementation of Epic than they did in LMR. If a larger proportion of the Epic alerts are less clinically important, users may override a greater number of them appropriately, leading to an overall reduction in acceptance rate, even though users are still complying with clinically appropriate DDI alerts. This explanation probably plays a role in the overall acceptance rates, as the knowledge base used in LMR is smaller than that used in our implementation of Epic. However, given that acceptance rates also fell for identical drug pairs from LMR to Epic, including pairs in the most severe tier that an expert panel had deemed “never” interactions, this does not fully explain our observed results.

A second possibility is that the increased number of alerts may reduce acceptance rates by causing alert fatigue.^{17, 32} Under one common conceptualization of alert fatigue, a high volume of DDI alerts distracts users from noticing more clinically important DDIs in a frequency-dependent fashion: as the number of less important alerts increases, the odds of missing a more important alert should increase. We explored this possibility with our natural experiment and found that reducing the interruptive DDI alert burden by more than half resulted in a modest increase in acceptance rates for the remaining alerts, lending support to this formulation of alert fatigue. Under another definition of alert fatigue, in addition to simply distracting users with noise, the increased number of alerts may also negatively affect users’ perceptions of the reliability of the alerting system. As mistrust in the validity of alerts increases, users may begin to devote less attention to alerts or ignore them entirely. Once credibility is lost, users may disregard even the most important alerts.^{33–35} In this scenario, decreasing alert burden after credibility has already been lost may not be enough to regain users’ trust. Evaluating this hypothesis will require further study.

A third potential reason is that other aspects of the interface or workflow affect alert acceptance. There were many differences in workflow between the two systems:

1. LMR alerts behaved differently based on tier. Tier 1 alerts were hard stops, Tier 2 alerts were interruptive with required override reason, and Tier 3 alerts were non-interruptive. This tiering was designed to focus users on the most severe DDI alerts (which they could not override), while deemphasizing the least severe. In our implementation of Epic, all three tiers are interruptive alerts with no override reason required. Although it is possible to configure Epic to require an override reason,

our organization chose not to do so. It is currently not possible to configure hard stops or non-interruptive DDI alerts in Epic.

2. LMR displays DDI alerts at the moment a medication is selected, while Epic displays them much later, at the time of signing. At this point, providers may be more “committed” to the drug, given that they completed considerable additional work. Thus, it is possible that showing alerts at the time of order signing increases the override rate versus showing alerts at the time of medication selection.
3. LMR required the user to accept or override each DDI alert on a pair-by-pair basis. In our implementation of Epic, when multiple alerts are displayed at one time, the user can override all of them with a single click. Further, since providers can sign medication orders in batch in Epic, alerts for many medications may appear at the same time and thus be overridden all together.

Based on results from a prior study of tiered DDI alerts in LMR,³⁶ we believe that workflow differences (particularly the difference in tiering) likely play the most significant role in the drop in acceptance rates. The first version of LMR, like Epic, presented all three tiers of DDIs interruptively and required an override reason for all tiers. When tiering was added to LMR, acceptance rates increased from 34 to 100% for Tier 1 alerts (due to the hard stop that prevented co-prescription of both drugs) and from 10 to 29% for Tier 2 alerts.³⁶ This dramatic change in acceptance rates after tiering is consistent with our experience in Epic, which moved in the opposite direction, strongly supporting the hypothesis that the lack of differentiation among tiers in Epic is a major driver of the overall drop in alert acceptance rates.

Implications

Our study has several important implications. First, it suggests that changing from a highly tailored DDI alerting system to a more general one as part of an EHR conversion can decrease the acceptance rate of DDI alerts and increase alert burden on users. Second, it provides a quantification of the impact of alert fatigue. Though alert fatigue has been described frequently in the literature, it has not previously been quantified experimentally. Our data from before and after Tier 3 suppression in Epic suggests that the effects of alert fatigue are real but, in our case, outweighed by other differences.

Strengths and Limitations

A key strength of our study is the before-after design at one institution, suggesting that the observed changes in alert burden and acceptance rates were due to the change in the EHR rather than population change. Further, the natural experiment and prior data comparing tiered and non-tiered alerting allow for a comprehensive analysis of potential underlying causes of these changes, including a novel quantification of the effect of

alert fatigue. Our study also has several limitations. First, it was done only at a single site, and with a single commercial EHR—further study with other systems would be informative. Second, it looks at only the initial 1-year period after implementation of the new commercial EHR system; alert acceptance may still increase or decrease after a longer period of stabilization (although as of 2-year post-implementation, there is no trend in the data to suggest that this will spontaneously happen). Third, the study does not investigate ADE rates before and after the change in system and thus does not provide a direct measure of the change on patient outcomes. Fourth, there were a number of differences between the EHR user interfaces, and we did not attempt to study each feature separately. A randomized trial of different alerting features (including tiering, workflow, user interface, and knowledge base) would be the gold standard for determining the relative contribution of each of these features. Finally, there were artifacts in the data we were unable to explain, such as changes in the number of alerts that fired for a single drug pair. Despite these limitations, the increase in alert burden and reduction in alert acceptance is quite substantial.

CONCLUSION

In summary, understanding the effects of alert design and alert fatigue for DDI and other medication alerts is crucial to implementing effective alerts that reduce ADEs and make patients safer. We found that switching from a homegrown EHR with strongly tiered interruptive DDI alerting driven by a highly customized knowledge base to a widely used commercial EHR with weakly tiered interruptive DDI alerting driven by a third-party knowledge base without customization resulted in a significant increase in alert burden and decrease in alert acceptance rates. The decrease in acceptance rates cannot be fully explained by the difference between the drug pairs present in each knowledge base since the decrease in acceptance rates is also observed for matched individual drug pairs. Furthermore, the decrease in acceptance rates cannot be neatly explained simply by the alert fatigue associated with the increased alert burden in the new system. Disabling more than 50% of the DDI alerts in the new system only marginally improved acceptance rates. Based on a prior study which showed much larger acceptance rate increases after introducing carefully delineated severity tiers,³⁶ the loss of strongly differentiated tiering in the new system may also be a driving factor behind the drop in acceptance rates. It is likely that each of these factors plays a role. Organizations creating changes to their DDI alerting interface should monitor the impact on alert burden and acceptance rates.

Contributors: None other than the authors as listed.

Corresponding Author: Adam Wright, PhD; Division of General Internal Medicine and Primary Care Brigham and Women's Hospital and Harvard Medical School, Boston, MA, USA (e-mail: awright@bwh.harvard.edu).

Funding Information Research reported in this publication was supported by the National Library of Medicine (NLM) of the National Institutes of Health under Award Number R01LM011966 and the Agency for Healthcare Research and Quality (AHRQ) Center for Education and Research on Therapeutics grant 1U18HS016970. The content is solely the responsibility of the authors and does not necessarily represent the official views of NIH or AHRQ. Neither NIH nor AHRQ were involved in the design or execution of the project or the decision to publish the results.

Compliance with Ethical Standards:

IRB Approval: Approval for the study was obtained from the Partners HealthCare Institutional Review Board.

Prior Presentation: Early data from this study was presented as a poster at the 2016 AMIA Annual Symposium.

Conflict of Interest: The authors declare that they do not have a conflict of interest.

REFERENCES

1. **Classen DC, Pestotnik SL, Evans RS, Lloyd JF, Burke JP.** Adverse drug events in hospitalized patients. Excess length of stay, extra costs, and attributable mortality. *JAMA*. 1997;277(4):301–6.
2. **Dechanont S, Maphanta S, Butthum B, Kongkaew C.** Hospital admissions/visits associated with drug-drug interactions: a systematic review and meta-analysis. *Pharmacoepidemiol Drug Saf*. 2014;23(5):489–97.
3. **Juurink DN, Mamdani M, Kopp A, Laupacis A, Redelmeier DA.** Drug-drug interactions among elderly patients hospitalized for drug toxicity. *JAMA*. 2003;289(13):1652–8.
4. **Johnson JA, Bootman JL.** Drug-related morbidity and mortality. A cost-of-illness model. *Archives of internal medicine*. 1995;155(18):1949–56.
5. **Melmon KL.** Preventable drug reactions—causes and cures. *The New England journal of medicine*. 1971;284(24):1361–8.
6. **Bates DW, Cullen DJ, Laird N, et al.** Incidence of adverse drug events and potential adverse drug events. Implications for prevention. ADE Prevention Study Group. *JAMA*. 1995;274(1):29–34.
7. **Rothschild JM, Federico FA, Gandhi TK, et al.** Analysis of medication-related malpractice claims: causes, preventability, and costs. *Archives of internal medicine*. 2002;162(21):2414–20.
8. **Ammenwerth E, Schnell-Inderst P, Machan C, Siebert U.** The effect of electronic prescribing on medication errors and adverse drug events: a systematic review. *Journal of the American Medical Informatics Association*: JAMIA. 2008;15(5):585–600.
9. **Bates DW, Leape LL, Cullen DJ, et al.** Effect of computerized physician order entry and a team intervention on prevention of serious medication errors. *JAMA*. 1998;280(15):1311–6.
10. **Dumbreck S, Flynn A, Nairn M, et al.** Drug-disease and drug-drug interactions: systematic examination of recommendations in 12 UK national clinical guidelines. *BMJ*; 2015.
11. **Wright A, Hickman TT, McEvoy D, et al.** Analysis of clinical decision support system malfunctions: a case series and survey. *Journal of the American Medical Informatics Association* : JAMIA. 2016.
12. **Kaushal R, Jha AK, Franz C, et al.** Return on investment for a computerized physician order entry system. *Journal of the American Medical Informatics Association* : JAMIA. 2006;13(3):261–6.
13. **Kuperman GJ, Bates DW, Teich JM, Schneider JR, Cheiman D.** A new knowledge structure for drug-drug interactions. *Proceedings. Symposium on Computer Applications in Medical Care*. 1994:836–40.
14. **Bates DW, Leape LL, Petrycki S.** Incidence and preventability of adverse drug events in hospitalized adults. *Journal of General Internal Medicine*. 1993;8(6):289–94.
15. **Medicare and Medicaid Programs; Electronic Health Record Incentive Program—Stage 3 and Modifications to Meaningful Use in 2015 Through 2017**, 80 Fed. Reg 200 (October 16, 2015). *Federal Register*.
16. **Hansten PD, Horn JR.** Hansten and Horn Drug Interactions. 2016.
17. **Phansalkar S, van der Sijs H, Tucker AD, et al.** Drug-drug interactions that should be non-interruptive in order to reduce alert fatigue in electronic health records. *Journal of the American Medical Informatics Association* : JAMIA. 2013;20(3):489–93.

18. **Phansalkar S, Desai AA, Bell D, et al.** High-priority drug-drug interactions for use in electronic health records. *Journal of the American Medical Informatics Association : JAMIA*. 2012;19(5):735–43.
19. **Peters LB, Bahr N, Bodenreider O.** Evaluating drug-drug interaction information in NDF-RT and DrugBank. *Journal of biomedical semantics*. 2015;6:19.
20. **Shah NR, Seger AC, Seger DL, et al.** Improving acceptance of computerized prescribing alerts in ambulatory care. *Journal of the American Medical Informatics Association : JAMIA*. 2006;13(1):5–11.
21. **Slight SP, Seger DL, Nanji KC, et al.** Are we heeding the warning signs? Examining providers' overrides of computerized drug-drug interaction alerts in primary care. *PLoS One*. 2013;8(12).
22. **Horn JR, Hansten PD.** Another thiopurine-allopurinol tragedy. *Pharmacy Times*; 2015.
23. **Horn JR, Hansten PD.** More evidence of a warfarin-antibiotic interaction. *Pharmacy Times*. 2009;75(6).
24. **Niece KL, Boyd NK, Akers KS.** In vitro study of the variable effects of proton pump inhibitors on voriconazole. *Antimicrobial Agents and Chemotherapy*. 2015;59(9):5548–54.
25. **Westergren T, Johansson P, Molden E.** Probable warfarin-simvastatin interaction. *The Annals of Pharmacotherapy*. 2007;41(7):1292–5.
26. **Granfors MT, Backman JT, Neuvonen M, Neuvonen PJ.** Ciprofloxacin greatly increases concentrations and hypotensive effect of tizanidine by inhibiting its cytochrome P450 1A2-mediated presystemic metabolism. *Clinical Pharmacology and Therapeutics*. 2004;76(6):598–606.
27. **Methaneethorn J, Chaiwong K, Pongpanich K, Sonsingh P, Lohitnavy M.** A pharmacokinetic drug-drug interaction model of simvastatin and clarithromycin in humans. *Conference proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual Conference*. 2014;2014:5703–6.
28. **Patel AM, Shariff S, Bailey DG, et al.** Statin toxicity from macrolide antibiotic coprescription: a population-based cohort study. *Annals of internal medicine*. 2013;158(12):869–76.
29. **Lozano R, Bibian C, Guilez RM, et al.** Clinical relevance of the (S)-citalopram-omeprazole interaction in geriatric patients. *British Journal of Clinical Pharmacology*. 2014;77(6):1086–7.
30. **Nishio S, Watanabe H, Kosuge K, et al.** Interaction between amlodipine and simvastatin in patients with hypercholesterolemia and hypertension. *Hypertension research : official journal of the Japanese Society of Hypertension*. 2005;28(3):223–7.
31. **Lane MA, Zeringue A, McDonald JR.** Serious bleeding events due to warfarin and antibiotic co-prescription in a cohort of veterans. *The American Journal of Medicine*. 2014;127(7):657–63.e2.
32. **Phansalkar S, Desai A, Choksi A, et al.** Criteria for assessing high-priority drug-drug interactions for clinical decision support in electronic health records. *BMC medical informatics and decision making*. 2013;13(1):65.
33. **Peterson J.** Integrating multiple medication decision support systems: how will we make it all work? *AHRQ Perspectives on Safety: AHRQ*; 2008.
34. **van der Sijs H, Aarts J, Vulto A, Berg M.** Overriding of drug safety alerts in computerized physician order entry. *Journal of the American Medical Informatics Association: JAMIA*. 2006;13(2):138–47.
35. **Topaz M, Seger DL, Slight SP, et al.** Rising drug allergy alert overrides in electronic health records: an observational retrospective study of a decade of experience. *Journal of the American Medical Informatics Association : JAMIA*. 2016;23(3):601–8.
36. **Paterno MD, Maviglia SM, Gorman PN, et al.** Tiering drug-drug interaction alerts by severity increases compliance rates. *Journal of the American Medical Informatics Association : JAMIA*. 2009;16(1):40–6.